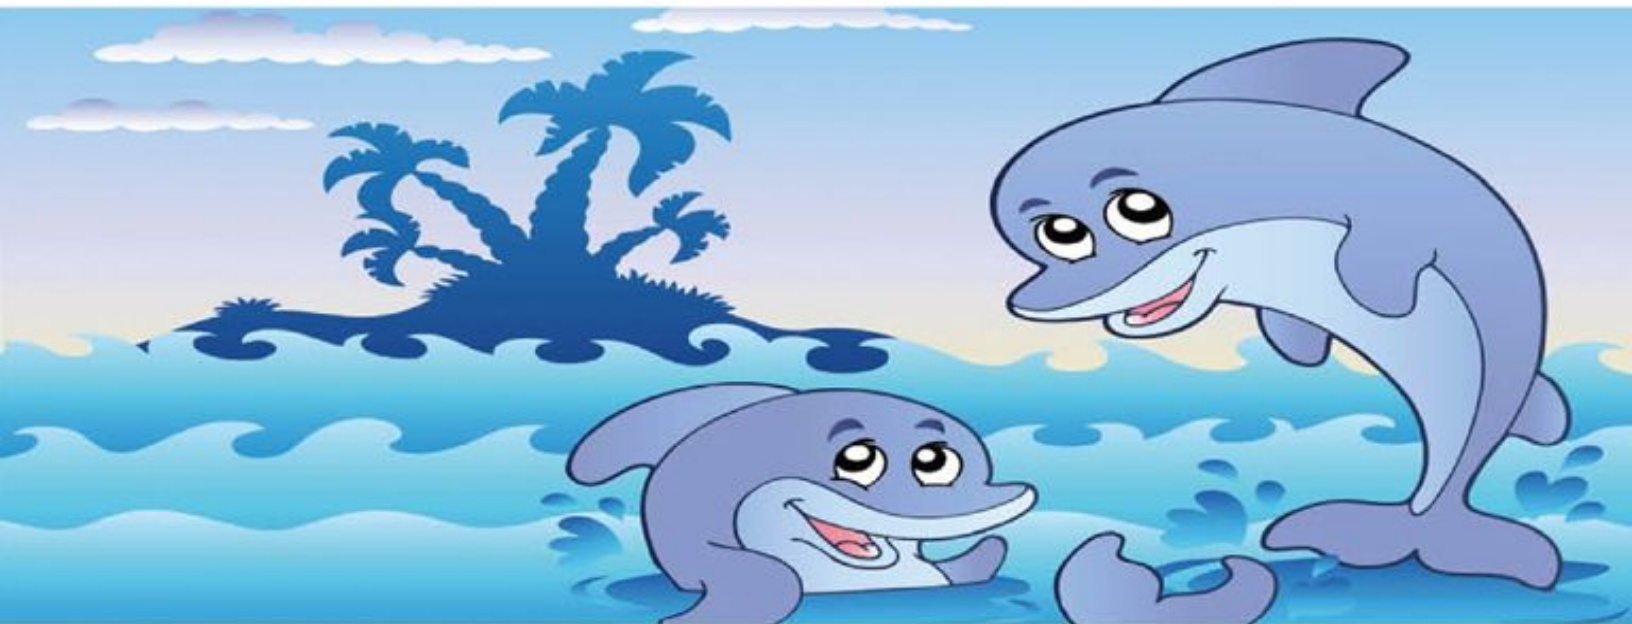


Inside MySQL: InnoDB Storage Engine, Second Edition

MySQL技术内幕

InnoDB存储引擎

第2版



姜承尧◎著



机械工业出版社
China Machine Press



[目次](#)

[序](#)

[第1章 数据库系统概论](#)

[1.1 数据库系统的发展](#)

[1.2 数据库系统的组成](#)

[1.3 数据库系统的层次结构](#)

[1.4 数据库系统的组成](#)

[1.5 数据库系统的组成](#)

[第2章 MySQL数据库系统](#)

[2.1 数据库系统的发展](#)

[2.2 MySQL数据库](#)

[2.3 MySQL数据库](#)

[2.3.1 InnoDB数据库](#)

[2.3.2 MyISAM数据库](#)

[2.3.3 NDB数据库](#)

[2.3.4 Memory数据库](#)

[2.3.5 Archive数据库](#)

[2.3.6 Federated数据库](#)

[1.3.7 MariaDB](#)

[1.3.8 参考](#)

[1.4 参考](#)

[1.5 MySQL](#)

[1.5.1 TCP/IP](#)

[1.5.2 参考](#)

[1.5.3 UNIX](#)

[1.6 参考](#)

[2 InnoDB](#)

[2.1 InnoDB](#)

[2.2 InnoDB](#)

[2.3 InnoDB](#)

[2.3.1 参考](#)

[2.3.2 参考](#)

[2.4 Checkpoint](#)

[2.5 Master Thread](#)

[2.5.1 InnoDB 1.0.x Master Thread](#)

[2.5.2 InnoDB 1.2.x Master Thread](#)

[2.5.3 InnoDB 1.2.x Master Thread](#)

[2.6 InnoDB](#)

[2.6.1 数据库](#)

[2.6.2 索引](#)

[2.6.3 数据库事务](#)

[2.6.4 数据库IO](#)

[2.6.5 数据库备份](#)

[2.7 数据库性能优化](#)

[2.8 数据库安全](#)

[3 数据库应用](#)

[3.1 数据库应用](#)

[3.1.1 数据库应用](#)

[3.1.2 数据库应用](#)

[3.2 数据库应用](#)

[3.2.1 数据库应用](#)

[3.2.2 数据库应用](#)

[3.2.3 数据库应用](#)

[3.2.4 数据库应用](#)

[3.3 数据库应用](#)

[3.4 pid管理](#)

[3.5 数据库应用](#)

[3.6 InnoDB数据库](#)

[3.6.1](#)

[3.6.2](#)

[3.7](#)

[4](#)

[4.1](#)

[4.2](#) [InnoDB](#)

[4.2.1](#)

[4.2.2](#)

[4.2.3](#)

[4.2.4](#)

[4.2.5](#)

[4.3](#) [InnoDB](#)

[4.3.1](#) [Compact](#)

[4.3.2](#) [Redundant](#)

[4.3.3](#)

[4.3.4](#) [Compressed](#)[Dynamic](#)

[4.3.5](#) [CHAR](#)

[4.4](#) [InnoDB](#)

[4.4.1](#) [File Header](#)

[4.4.2](#) [Page Header](#)

[4.4.3 Infimum□Supremum Record](#)

[4.4.4 User Record□Free Space](#)

[4.4.5 Page Directory](#)

[4.4.6 File Trailer](#)

[4.4.7 InnoDB□□□□□□□□](#)

[4.5 Named File Formats□□](#)

[4.6 □□](#)

[4.6.1 □□□□□](#)

[4.6.2 □□□□□□□□](#)

[4.6.3 □□□□□□□□](#)

[4.6.4 □□□□□□□□](#)

[4.6.5 ENUM□SET□□](#)

[4.6.6 □□□□□□](#)

[4.6.7 □□□□](#)

[4.7 □□](#)

[4.7.1 □□□□□](#)

[4.7.2 □□□□](#)

[4.8 □□□](#)

[4.8.1 □□□□](#)

[4.8.2 □□□□](#)

[4.8.3](#) [索引](#)

[4.8.4](#) [索引的NULL值](#)

[4.8.5](#) [索引的创建](#)

[4.8.6](#) [索引的删除与重建](#)

[4.9](#) [视图](#)

[第5章](#) [数据库系统](#)

[5.1](#) [InnoDB存储引擎](#)

[5.2](#) [数据库的备份与恢复](#)

[5.2.1](#) [数据库的备份](#)

[5.2.2](#) [数据库的恢复](#)

[5.3](#) [B+树](#)

[5.3.1](#) [B+树的结构](#)

[5.3.2](#) [B+树的遍历](#)

[5.4](#) [B+树的索引](#)

[5.4.1](#) [索引的创建](#)

[5.4.2](#) [索引的删除](#)

[5.4.3](#) [B+树的索引](#)

[5.4.4](#) [B+树的索引](#)

[5.5](#) [Cardinality](#)

[5.5.1](#) [索引的Cardinality](#)

[5.5.2 InnoDB Cardinality](#)

[5.6 B+](#)

[5.6.1 B+](#)

[5.6.2](#)

[5.6.3](#)

[5.6.4](#)

[5.6.5](#)

[5.6.6 Multi-Range Read](#)

[5.6.7 Index Condition Pushdown\(ICP\)](#)

[5.7](#)

[5.7.1](#)

[5.7.2 InnoDB](#)

[5.7.3](#)

[5.8](#)

[5.8.1](#)

[5.8.2](#)

[5.8.3 InnoDB](#)

[5.8.4](#)

[5.9](#)

[6](#)

[6.1 背景](#)

[6.2 lock&latch](#)

[6.3 InnoDB锁机制](#)

[6.3.1 锁类型](#)

[6.3.2 锁粒度](#)

[6.3.3 锁兼容性](#)

[6.3.4 锁等待](#)

[6.3.5 死锁](#)

[6.4 索引](#)

[6.4.1 索引3要素](#)

[6.4.2 索引Phantom Problem](#)

[6.5 事务](#)

[6.5.1 事务](#)

[6.5.2 事务隔离](#)

[6.5.3 事务一致性](#)

[6.6 视图](#)

[6.7 分区](#)

[6.7.1 分区类型](#)

[6.7.2 分区表](#)

[6.7.3 分区索引](#)

[6.8 备份](#)

[6.9 恢复](#)

[7 高可用](#)

[7.1 主从复制](#)

[7.1.1 原理](#)

[7.1.2 配置](#)

[7.2 主从复制](#)

[7.2.1 redo](#)

[7.2.2 undo](#)

[7.2.3 purge](#)

[7.2.4 group commit](#)

[7.3 主从复制](#)

[7.4 主从复制SQL日志](#)

[7.5 主从复制日志](#)

[7.6 主从复制](#)

[7.7 主从复制](#)

[7.7.1 MySQL主从复制](#)

[7.7.2 主从XA](#)

[7.8 主从复制](#)

[7.8.1 主从复制](#)

[7.8.2 备份数据](#)

[7.8.3 恢复数据](#)

[7.9 索引](#)

[7.10 用户](#)

[第8章 数据库管理](#)

[8.1 数据库概述](#)

[8.2 用户](#)

[8.3 备份](#)

[8.3.1 mysqldump](#)

[8.3.2 SELECT...INTO OUTFILE](#)

[8.3.3 物理备份](#)

[8.3.4 LOAD DATA INFILE](#)

[8.3.5 mysqlimport](#)

[8.4 数据库性能优化](#)

[8.5 备份](#)

[8.5.1 ibbackup](#)

[8.5.2 XtraBackup](#)

[8.5.3 XtraBackup物理备份](#)

[8.6 索引](#)

[8.7 用户](#)

[8.7.1 配置数据库](#)

[8.7.2 配置+数据库配置](#)

[8.8 配置](#)

[9 配置](#)

[9.1 配置CPU](#)

[9.2 配置](#)

[9.3 配置数据库配置](#)

[9.3.1 配置](#)

[9.3.2 配置](#)

[9.4 配置RAID](#)

[9.4.1 RAID配置](#)

[9.4.2 RAID Write Back配置](#)

[9.4.3 RAID配置](#)

[9.5 配置](#)

[9.6 配置数据库配置](#)

[9.7 配置数据库配置](#)

[9.7.1 sysbench](#)

[9.7.2 mysql-tpcc](#)

[9.8 配置](#)

[10 InnoDB配置数据库配置](#)

[10.1 编译InnoDB插件](#)

[10.2 InnoDB插件](#)

[10.3 MySQL 5.1编译InnoDB插件](#)

[10.3.1 Windows](#)

[10.3.2 Linux](#)

[10.4 cmake编译InnoDB插件](#)

[10.5 编译](#)



It's fair to say that MySQL is the most popular open source database. It has a very large installed base and number of users. Let's see what are the reasons MySQL is so popular, where it stands currently, and maybe touch on some of its future (although predicting the future is rarely successful).

Looking at the customer area of MySQL, which includes Facebook, Flickr, Adobe (in Creative Suite 3), Drupal, Digg, LinkedIn, Wikipedia, eBay, YouTube, Google AdSense (source <http://mysql.com/customers/> and public resources), it's obvious that MySQL is everywhere. When you log in to your popular forum (powered by Bulleting) or blog (powered by WordPress), most likely it has MySQL as its backend database. Traditionally, two MySQL's characteristics, simplicity of use and performance, were what allowed it to gain such popularity. In addition to that, availability on a very wide range of platforms (including Windows) and built-in replication, which provides an easy scale-out solution for read-only clients, gave more user attractions and production deployments. There is simple evidence of MySQL's simplicity: In 15 minutes or less, you really can get installed, have a working database, and start running queries and store data. From its early stages MySQL had a good interface to most popular languages for Web development—PHP and Perl, and also Java and ODBC connectors.

There are two best known storage engines in MySQL: MyISAM and InnoDB (I don't cover NDB cluster here; it's a totally different story). MyISAM comes as the default storage engine and historically it is the oldest, but InnoDB is ACID compliant and provides transactions, row-

level locking, MVCC, automatic recovery and data corruption detection. This makes it the storage engine you want to choose for your application. Also, there is the third-party transaction storage engine PBXT, with characteristics similar to InnoDB, which is included in the MariaDB distribution.

MySQL's simplicity has its own drawback. Just as it is very easy to start working with it, it is very easy to start getting into trouble with it. As soon as your website or forum gets popular, you may figure out that the database is a bottleneck, and that you need special skills and tools to fix it.

The author of this book is a MySQL expert, especially in InnoDB storage engine. Hence, I highly recommend this book to new users of InnoDB as well as users who already have well-tuned InnoDB-based applications but need to get internal out of them.

Vadim Tkachenko

MySQL Performance Blog Percona CTO

MySQL MySQLPerformanceBlog.com

MySQL 2



資料庫設計

MySQL 資料庫設計原則
MySQL 資料庫設計原則
MySQL 資料庫設計原則
MySQL 資料庫設計原則

MySQL 資料庫設計原則
MySQL 資料庫設計原則
MySQL 資料庫設計原則
MySQL 資料庫設計原則

InnoDB 資料庫設計原則
Think Different
Think Different
20
90
Think Differently
Think Different
Different

DBA 資料庫設計原則
Why
What
Why
What
Why
What

資料庫設計原則

資料庫設計原則

資料庫設計原則

資料庫設計原則

資料庫設計原則

1 2

2 MySQL 5.6 InnoDB

1 2 50%

MySQL 5.6 InnoDB MySQL 5.6 InnoDB

InnoDB redo undo InnoDB undo InnoDB undo DBA InnoDB

6 InnoDB next-key locking InnoDB

InnoDB insert buffer insert buffer — change buffer

目錄

第一章 MySQL 簡介
1.1 MySQL InnoDB 引擎
1.2 MySQL DBA 工具
1.3 MySQL Replication
1.4 InnoDB 引擎

第二章 InnoDB 引擎
2.1 InnoDB 引擎的數據結構
2.2 MySQL 的數據庫引擎“引擎”

第三章 數據庫引擎

第四章 SQL

第五章 MySQL 引擎

第六章 C++、Python、Java

第七章 InnoDB 引擎
第八章 數據庫引擎

--	--	--	--	--	--

10 “ ”
10 InnoDB 3
3

```

00000000000000000000000000000000InnoDB0000000000
00000000000000000000API000000InnoDB000000000000
0000000000

```

www.hzbook.com

□□□□□

□□□□InnoDB□□□□□□□□□□□□□□□□□□□□□□□□□□□□
□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□
□□□□□□□□□□jiangchengyao@gmail.com□□□□□□□□□□
@insidemysql□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□
□□!



Pecona CEO Peter
Zaitsev CTO Vadim Tkachenko
InnoDB
InnoDB

1 MySQL

MySQL Linux Solaris
FreeBSD Mac Windows
MySQL MySQL
MySQL

3578 MySQL mysqld_safe
MySQL

MySQL Oracle
spfile Oracle
MySQL MySQL
MySQL

```
[root@xen-server bin]#mysql--help|grep my.cnf
```

```
order of preference,my.cnf,$MYSQL_TCP_PORT,
```

```
/etc/my.cnf/etc/mysql/my.cnf/usr/local/mysql/etc/my.cnf/.my.cnf
```

MySQL
/etc/my.cnf→/etc/mysql/my.cnf→usr/local/mysql/etc/my.cnf
→/.my.cnf“
MySQL”MySQL
Linux/etc/my.cnfWindows
.cnf.iniWindowsmysql--help

Default options are read from the following files in the given order:

C:\Windows\my.ini C:\Windows\my.cnf C:\my.ini C:\my.cnf C:\Program Files\MySQL\M

\MySQL Server 5.1\my.cnf

datadirLinuxdatadir
/usr/local/mysql/data

```
mysqlSHOW VARIABLES LIKE'datadir'\G;
```

```
*****1.row*****
```

Variable_name:datadir

Value:/usr/local/mysql/data/

1 row in set(0.00 sec)1 row in set(0.00 sec)

mysqlsystem ls-lh/usr/local/mysql/data

total 32K

drwxr-xr-x 2 root mysql 4.0K Aug 6 16:23 bin

drwxr-xr-x 2 root mysql 4.0K Aug 6 16:23 docs

drwxr-xr-x 3 root mysql 4.0K Aug 6 16:04 include

drwxr-xr-x 3 root mysql 4.0K Aug 6 16:04 lib

drwxr-xr-x 2 root mysql 4.0K Aug 6 16:23 libexec

drwxr-xr-x 10 root mysql 4.0K Aug 6 16:23 mysql-test

drwxr-xr-x 5 root mysql 4.0K Aug 6 16:04 share

drwxr-xr-x 5 root mysql 4.0K Aug 6 16:23 sql-bench

lrwxrwxrwx 1 root mysql 16 Aug 6 16:05 data- /opt/mysql_data/

data /opt/mysql_data
/opt/mysql_datamysqlMySQL
mysql:mysql

1.2 MySQL

DBA MySQL MySQL SQL Server Oracle DB2

MySQL

MySQL

1000 MySQL

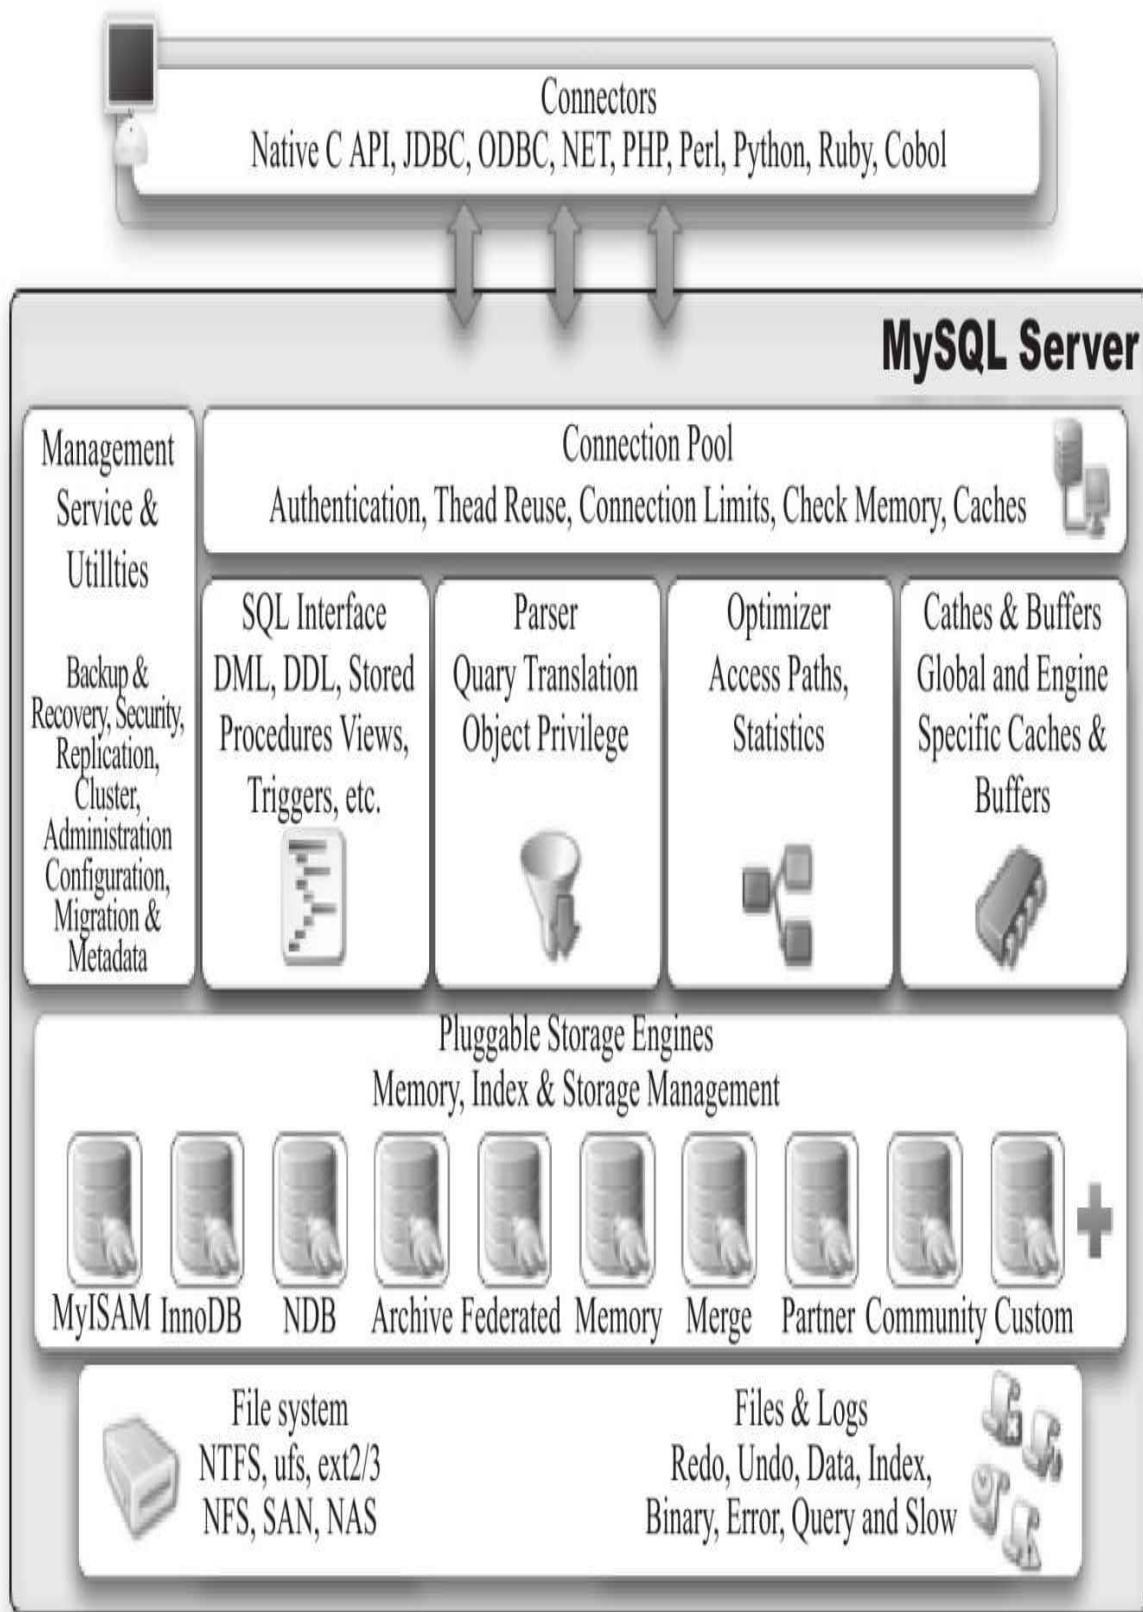
.....

MySQL

MySQL MySQL Oracle Microsoft SQL Server MySQL

SELECT INSERT UPDATE DELETE Oracle SQL Server MySQL

MySQL 1-1 MySQL



1-1 MySQL

1-1 MySQL

MySQL

MySQL

SQL

MySQL

MySQL

MySQL Cache

MySQL

MySQL

1-1 MySQL
MySQL
SQL
MySQL

1-1 MySQL
MySQL

1.3 MySQL

1.2 MySQL MySQL DBA MySQL

MySQL MySQL
eBay Igor Chernyshev MySQL Memory
<http://code.google.com/p/mysql-heap-dynamic-rows/>
eBay Personalization Platform Google Facebook
InnoDB SSD [1]
SSD MySQL
MySQL 16

MySQL MySQL InnoDB Oracle MySQL OLTP Online Transaction Processing

1.3.1 InnoDB

InnoDBOLTP
OracleMySQL5.5.8
InnoDB

```
InnoDB[REDACTED]InnoDB[REDACTED]
[REDACTED]MySQL 4.1[REDACTED]InnoDB[REDACTED]
[REDACTED]ibd[REDACTED]InnoDB[REDACTED]row disk[REDACTED]
```

InnoDB MVCC SQL 4 REPEATABLE next-key locking

phantom InnoDB insert
buffer double write adaptive hash index
read ahead

InnoDB clustered
InnoDB 6
ROWID

InnoDB MySQL Facebook Google
Yahoo InnoDB

[1] http://code.google.com/p/david-mysql-tools/wiki/innodb_secondary_buffer_pool

1.3.2 MyISAM

MyISAM
MySQL 5.5.8
MyISAM
Windows
MyISAM
ETL
MyISAM
cache

```
MyISAM MYD MYI MYD MYI
myisampack myisampack
Huffman myisampack
myisampack
```

```
MySQL 5.0      MyISAM      4GB      4GB
MyISAM      MAX_ROWS      AVG_ROW_LENGTH
MySQL 5.0      MyISAM      256TB
```

MySQL

MyISAM MySQL
LRU MySQL 5.1.23
32 64 4GB
64 4GB

1.3.3 NDB

2003 MySQL AB Sony Ericsson NDB 1-1
NDB Oracle RAC Oracle RAC
share everything share nothing
NDB MySQL 5.1
primary key lookups NDB
Data Node

NDB JOIN MySQL
NDB

MySQL NDB Cluster NDB Cluster
Carrier Grade Edition
<http://dev.mysql.com/downloads/cluster/index.html>
NDB Cluster

1.3.4 Memory

Memory is used for HEAP tables and for temporary tables. Memory is also used for the B+ tree.

Memory is used for TEXT, BLOB, varchar, and char. Memory is also used for eBay and Igor Chernyshev. Memory is also used for patch.

MySQL uses Memory for intermediate result. Memory is also used for TEXT, BLOB, MySQL, and MyISAM. Memory is also used for MyISAM.

1.3.5 Archive

Archive INSERT SELECT MySQL 5.1
Archive zlib row 1:10
Archive Archive
Archive

1.3.6 Federated

Federated MySQL
SQL Server Oracle Federated
MySQL

1.3.7 MariaDB

MariaDB is a community-developed database that is a fork of MySQL. It is designed to be a drop-in replacement for MySQL, but with several improvements. MariaDB is developed by MariaDB Corporation, which was founded by Michael Widenius, the creator of MySQL. MariaDB is known for its performance, reliability, and security. It supports many of the same features as MySQL, but also has some unique features, such as the MVCC (Multi-Version Concurrency Control) and BLOB (Binary Large Object) types.

1.3.8 数据库

数据库7数据库MySQL数据库数据库数据库Merge CSV
Sphinx Infobright数据库数据库数据库数据库MySQL数据库
数据库数据库数据库1.2数据库数据库

MySQL数据库数据库MySQL数据库MyISAM InnoDB 1.2
Sphinx数据库数据库数据库

MySQL数据库数据库数据库MySQL数据库MyISAM数据库数据库
InnoDB数据库“”数据库数据库数据库ETL数据库MyISAM数据库数据库
OLTP数据库InnoDB数据库数据库数据库

数据库数据库1000数据库MySQL数据库数据库MySQL数据库数据库
数据库数据库数据库数据库数据库数据库数据库数据库数据库数据库数据库
数据库数据库MySQL数据库数据库数据库Mytrix Inc. InnoDB数据库
1 TB数据库数据库数据库InnoDB数据库数据库/数据库数据库800/数据库

1.4 数据库事务

在1.3节中，我们介绍了MySQL的存储引擎。在1-2节中，我们介绍了MySQL的索引。在本节中，我们将介绍MySQL的事务。事务是数据库中的一个操作序列，这些操作要么都执行，要么都不执行。事务的四个特性是：原子性、一致性、隔离性和持久性。MySQL支持多种事务隔离级别，其中MVCC（多版本并发控制）是InnoDB存储引擎默认的事务隔离级别。

Feature	MyISAM	BDB	Memory	InnoDB	Archive	NDB
Storage Limits	No	No	Yes	64TB	No	Yes
Transactions (commit, rollback, etc.)		✓		✓		
Locking granularity	Table	Page	Table	Row	Row	Row
MVCC/Snapshot Read				✓	✓	✓
Geospatial support	✓					
B-Tree indexes	✓	✓	✓	✓		✓
Hash indexes			✓	✓		✓
Full text search index	✓					
Clustered index				✓		
Data Caches			✓	✓		✓
Index Caches	✓		✓	✓		✓
Compressed data	✓				✓	
Encrypted data (via function)	✓	✓	✓	✓	✓	✓
Storage cost (space used)	Low	Low	N/A	High	Very Low	Low
Memory cost	Low	Low	Medium	High	Low	High
Bulk Insert Speed	High	High	High	Low	Very High	High
Cluster database support						✓
Replication support	✓	✓	✓	✓	✓	✓
Foreign key support				✓		
Backup/Point-in-time recovery	✓	✓	✓	✓	✓	✓
Query cache support	✓	✓	✓	✓	✓	✓
Update Statistics for Data Dictionary	✓	✓	✓	✓	✓	✓

1-2 MySQL

MySQL은 다양한 엔진을 지원합니다. 이 문서에서는 MySQL의 기본 엔진인 InnoDB, MyISAM, 그리고 BLACKHOLE 엔진에 대해 설명합니다.

SHOW ENGINES 명령을 사용하여 MySQL의 엔진 정보를 조회할 수 있습니다. 이 명령은 information_schema.ENGINES 테이블에서 정보를 가져옵니다.

```
mysql> SHOW ENGINES\G;
```

```
*****1.row*****
```

```
Engine:InnoDB
```

```
Support:YES
```

```
Comment:Supports transactions,row-level locking,and foreign keys
```

```
Transactions:YES
```

```
XA:YES
```

```
Savepoints:YES
```

```
*****2.row*****
```

```
Engine:MRG_MYISAM
```

```
Support:YES
```

```
Comment:Collection of identical MyISAM tables
```

```
Transactions:NO
```

```
XA:NO
```

```
Savepoints:NO
```

```
*****3.row*****
```

```
Engine:BLACKHOLE
```

```
Support:YES
```

```
Comment:/dev/null storage engine(anything you write to it disappears)
```

```
Transactions:NO
```

```
XA:NO
```

```
Savepoints:NO
```

*****4.row*****

Engine:CSV

Support:YES

Comment:CSV storage engine

Transactions:NO

XA:NO

Savepoints:NO

*****5.row*****

Engine:MEMORY

Support:YES

Comment:Hash based,stored in memory,useful for temporary tables

Transactions:NO

XA:NO

Savepoints:NO

*****6.row*****

Engine:FEDERATED

Support:NO

Comment:Federated MySQL storage engine

Transactions:NULL

XA:NULL

Savepoints:NULL

*****7.row*****

Engine:ARCHIVE

Support:YES

Comment:Archive storage engine

Transactions:NO

XA:NO

Savepoints:NO

*****8.row*****

Engine:MyISAM

```

MySQLMySQLSQL Server
AdventureWorksOracleMySQL
OracleSQL
Server
http://dev.mysql.com/doc/

```


1.5 MySQL

```

MySQL
MySQL
MySQL
TCP/IP
UNIX
MySQL

```

1.5.1 TCP/IP

TCP/IP MySQL
client
server
TCP/IP
Windows Linux MySQL

```
C:\>mysql-h192.168.0.101-u david-p
```

Enter password:

Welcome to the MySQL monitor. Commands end with ; or \g.

Your MySQL connection id is 18358

Server version:5.0.77-log MySQL Community Server(GPL)

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

mysql

Windows Host IP 192.168.0.101 MySQL
TCP/IP MySQL DDL
DML

```

000000000000TCP/IP0000MySQL0000MySQL0000000000000000
000000000000IP00000000MySQL00000000mysql00000000user000
000

```

```
mysql>USE mysql;
```

Database changed

```
mysql>SELECT host,user,password FROM user;
```

```
*****1.row*****
```

```
host:192.168.24.%
```

```
user:root
```

```
password:*75DBD4FA548120B54FE693006C41AA9A16DE8FBE
```

```
*****2.row*****
```

```
host:nineyou0-43
```

```
user:root
```

```
password:*75DBD4FA548120B54FE693006C41AA9A16DE8FBE
```

```
*****3.row*****
```

```
host:127.0.0.1
```

```
user:root
```

```
password:*75DBD4FA548120B54FE693006C41AA9A16DE8FBE
```

```
*****4.row*****
```

```
host:192.168.0.100
```

```
user:zlm
```

```
password:*DAE0939275CC7CD8E0293812A31735DA9CF0953C
```

```
*****5.row*****
```

```
host:%
```

```
user:david
```

```
password:
```

```
5 rows in set(0.00 sec)
```

MySQL david IP
root

1.5.2 〇〇〇〇〇〇〇〇〇〇

[illegible]

1.5.3 UNIX

```

LinuxUNIXUNIXUNIXUNIX
MySQL
--socket=/tmp/mysql.sock
UNIX

```

```
mysql> SHOW VARIABLES LIKE 'socket';
```

```
*****1.row*****
```

Variable_name:socket

Value: /tmp/mysql.sock

```
1 row in set(0.00 sec)
```

UNIX

```
[root@stargazer ~]#mysql-udavid-S/tmp/mysql.sock
```

Welcome to the MySQL monitor. Commands end with ; or \g.

Your MySQL connection id is 20333

Server version:5.0.77-log MySQL Community Server(GPL)

```
Type 'help;' or '\h' for help.Type '\c' to clear the buffer.
```

mysql

1.6 概要

本書は、MySQLのインストール、設定、運用、保守に関する情報を提供することを目的としています。MySQL DBAの役割、MySQLのアーキテクチャ、MySQLのインストール、MySQLのバックアップと復元、MySQLのセキュリティ、MySQLの性能チューニング、MySQLのトラブルシューティングについて詳しく説明します。

本書は、MySQLのインストール、設定、運用、保守に関する情報を提供することを目的としています。MySQLのアーキテクチャ、MySQLのインストール、MySQLのバックアップと復元、MySQLのセキュリティ、MySQLの性能チューニング、MySQLのトラブルシューティングについて詳しく説明します。

第2章 InnoDB引擎

InnoDB引擎是MySQL数据库的默认存储引擎，它也是Oracle数据库的默认存储引擎。InnoDB引擎是OLTP（Online Transaction Processing）数据库的理想选择，而MySQL数据库的默认存储引擎是InnoDB引擎。InnoDB引擎是MySQL数据库的默认存储引擎，也是Oracle数据库的默认存储引擎。InnoDB引擎是OLTP（Online Transaction Processing）数据库的理想选择，而MySQL数据库的默认存储引擎是InnoDB引擎。InnoDB引擎是MySQL数据库的默认存储引擎，也是Oracle数据库的默认存储引擎。InnoDB引擎是OLTP（Online Transaction Processing）数据库的理想选择，而MySQL数据库的默认存储引擎是InnoDB引擎。

2.1 InnoDB引擎

InnoDB是MySQL数据库的默认存储引擎。它是由MySQL的创始人Heikki Tuuri在1994年创建的。InnoDB引擎支持事务（ACID）、行级锁、崩溃恢复、MVCC（多版本并发控制）以及外键。它比MySQL的默认存储引擎MyISAM更强大，但CPU使用率更高。

Heikki Tuuri在1964年出生。他在1990年创建了InnoDB引擎。Linus Torvalds在1990年创建了Linux操作系统。Innobase Oy是InnoDB的原始开发公司，其CEO是Heikki Tuuri。Calvin Sun和Jimmy Yang也是InnoDB的早期开发人员。Sybase是另一个数据库公司。

InnoDB引擎被广泛应用于Google、Yahoo!、Facebook、YouTube、Flickr、Second Life、MySQL、InnoDB OLTP、MySQL InnoDB等。

MySQL数据库在互联网上非常流行。Slashdot.org是MySQL的早期用户之一。InnoDB引擎由Myrix Inc.开发，支持1 TB的数据库。InnoDB引擎/数据库支持800个并发连接。InnoDB引擎支持多版本并发控制（MVCC）。

InnoDB引擎是MySQL数据库的默认存储引擎。它遵循GNU GPL 2许可证。MySQL数据库的官方网站是<http://www.mysql.com/about/legal/>。

[1] 2006年，Oracle公司收购了MySQL数据库。

2.2 InnoDB

InnoDB MySQL MySQL
MySQL 5.1 MySQL
MySQL MySQL 5.1
InnoDB InnoDB InnoDB
InnoDB InnoDB Plugin InnoDB 1.0.x MySQL
5.5 InnoDB 1.1.x MySQL 5.6 InnoDB
1.2.x 2-1 InnoDB

表 2-1 InnoDB 各版本功能对比

版 本	功 能
老版本 InnoDB	支持 ACID、行锁设计、MVCC
InnoDB 1.0.x	继承了上述版本所有功能，增加了 compress 和 dynamic 页格式
InnoDB 1.1.x	继承了上述版本所有功能，增加了 Linux AIO、多回滚段
InnoDB 1.2.x	继承了上述版本所有功能，增加了全文索引支持、在线索引添加

MySQL MySQL 5.1 InnoDB
Plugin DBA InnoDB Plugin InnoDB 1.1
2-1 compress dynamic InnoDB
Plugin Linux Native AIO InnoDB

Plugin1023MySQL 5.5
InnoDB Plugin

2.3 InnoDB

1 MySQL InnoDB
2-1 InnoDB InnoDB
.....

□ /

□

□ redo log

.....

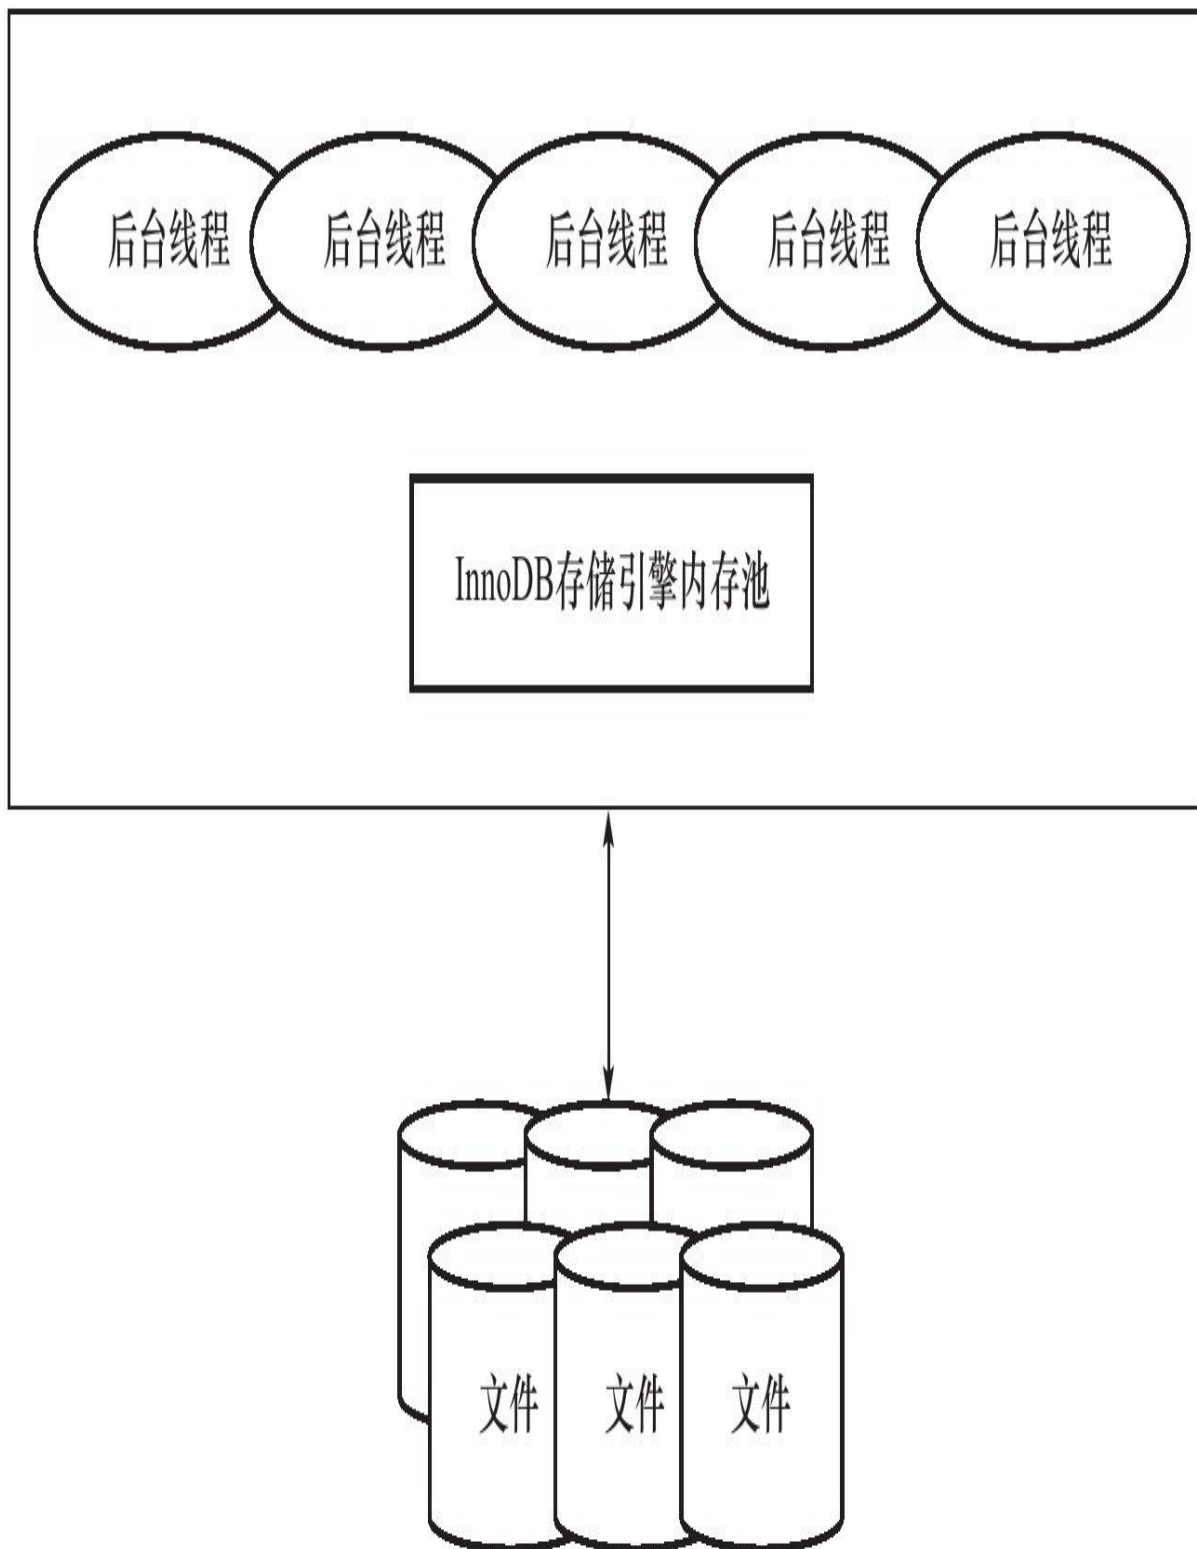


图 2-1 InnoDB架构图

MySQL InnoDB 引擎的线程池结构

2.3.1 线程池

InnoDB 线程池的线程池结构

1.Master Thread

Master Thread 是 InnoDB 线程池中的主线程，负责管理整个线程池。它负责分配和回收线程，并负责管理 INSERT BUFFER 和 UNDO 空间。在 MySQL 2.5 版本中，Master Thread 负责管理整个线程池。

2.IO Thread

InnoDB 线程池中的 IO 线程分为两种：Async IO 线程和 Sync IO 线程。Async IO 线程负责处理异步 IO 请求，而 Sync IO 线程负责处理同步 IO 请求。在 InnoDB 1.0 版本中，IO 线程池的大小为 4。在 Linux 系统上，IO 线程池的大小可以通过 `innodb_file_io_threads` 参数来设置。在 Windows 系统上，IO 线程池的大小可以通过 `innodb_read_io_threads` 和 `innodb_write_io_threads` 参数来设置。在 InnoDB 1.0.x 版本中，IO 线程池的大小可以通过 `innodb_read_io_threads` 和 `innodb_write_io_threads` 参数来设置。

```
mysql>SHOW VARIABLES LIKE'innodb_version'\G;

*****1.row*****

Variable_name:innodb_version

Value:1.0.6

1 row in set(0.00 sec)

mysql>SHOW VARIABLES LIKE'innodb_%io_threads'\G;

*****1.row*****

Variable_name:innodb_read_io_threads

Value:4

*****2.row*****
```

Variable_name:innodb_write_io_threads

Value:4

2 rows in set(0.00 sec)

SHOW ENGINE INNODB STATUSInnoDBIO Thread

mysql>SHOW ENGINE INNODB STATUS\G;

*****1.row*****

Type:InnoDB

Name:

Status:

=====

100719 21:55:26 INNODB MONITOR OUTPUT

=====

Per second averages calculated from the last 36 seconds

.....

FILE I/O

I/O thread 0 state:waiting for i/o request(insert buffer thread)

I/O thread 1 state:waiting for i/o request(log thread)

I/O thread 2 state:waiting for i/o request(read thread)

I/O thread 3 state:waiting for i/o request(read thread)

I/O thread 4 state:waiting for i/o request(read thread)

I/O thread 5 state:waiting for i/o request(read thread)

I/O thread 6 state:waiting for i/o request(write thread)

I/O thread 7 state:waiting for i/o request(write thread)

I/O thread 8 state:waiting for i/o request(write thread)

I/O thread 9 state:waiting for i/o request(write thread)

.....

END OF INNODB MONITOR OUTPUT

=====

1 row in set (0.01 sec)

IO Thread 0 insert buffer thread IO Thread 1 log
thread innodb_read_io_threads
innodb_write_io_threads ID

3.Purge Thread

undolog PurgeThread
undo InnoDB 1.1 purge InnoDB Master
Thread InnoDB 1.1 purge
Master Thread CPU MySQL
Purge Thread

[mysqld]

innodb_purge_threads=1

InnoDB 1.1 innodb_purge_threads=1 InnoDB
1

120529 22:54:16[Warning]option'innodb-purge-threads':unsigned value 4 adjusted to 1

InnoDB 1.2 InnoDB Purge Thread
undo Purge Thread undo
4 Purge Thread

mysql>SELECT VERSION()\G;

*****1.row*****

```
VERSION():5.6.6
```

```
1 row in set(0.00 sec)
```

```
mysql>SHOW VARIABLES LIKE'innodb_purge_threads'\G;
```

```
*****1.row*****
```

```
Variable_name:innodb_purge_threads
```

```
Value:4
```

```
1 row in set(0.00 sec)
```

4. Page Cleaner Thread

Page Cleaner Thread InnoDB 1.2.x
Master Thread
InnoDB

2.3.2 配置

1. 简介

InnoDB 是 MySQL 数据库的默认存储引擎。它是一个基于磁盘的数据库（Disk-based Database），它使用 CPU 缓存数据，并定期将数据写入磁盘。它支持事务、行级锁、崩溃恢复等功能。

InnoDB 数据库的默认配置参数如下所示。如果配置参数与默认值不同，则需要在配置文件中指定。如果配置参数与默认值相同，则不需要指定。如果配置参数与默认值不同，则需要使用“FIX”来固定配置参数。

InnoDB 数据库的默认配置参数如下所示。如果配置参数与默认值不同，则需要在配置文件中指定。如果配置参数与默认值相同，则不需要指定。如果配置参数与默认值不同，则需要使用“Checkpoint”来检查配置参数。

InnoDB 数据库的默认配置参数如下所示。如果配置参数与默认值不同，则需要在配置文件中指定。如果配置参数与默认值相同，则不需要指定。如果配置参数与默认值不同，则需要使用“32”来指定配置参数。如果配置参数与默认值不同，则需要使用“3G”来指定配置参数。如果配置参数与默认值不同，则需要使用“PAE”来指定配置参数。如果配置参数与默认值不同，则需要使用“32”来指定配置参数。如果配置参数与默认值不同，则需要使用“64GB”来指定配置参数。如果配置参数与默认值不同，则需要使用“8GB”来指定配置参数。如果配置参数与默认值不同，则需要使用“PC”来指定配置参数。如果配置参数与默认值不同，则需要使用“512GB”来指定配置参数。如果配置参数与默认值不同，则需要使用“64”来指定配置参数。

InnoDB 数据库的默认配置参数如下所示。如果配置参数与默认值不同，则需要在配置文件中指定。如果配置参数与默认值相同，则不需要指定。如果配置参数与默认值不同，则需要使用“innodb_buffer_pool_size”来指定配置参数。如果配置参数与默认值不同，则需要使用“MySQL”来指定配置参数。如果配置参数与默认值不同，则需要使用“InnoDB”来指定配置参数。如果配置参数与默认值不同，则需要使用“15GB”来指定配置参数。

```
mysql> SHOW VARIABLES LIKE 'innodb_buffer_pool_size'\G;
```

```
*****1. row*****
```

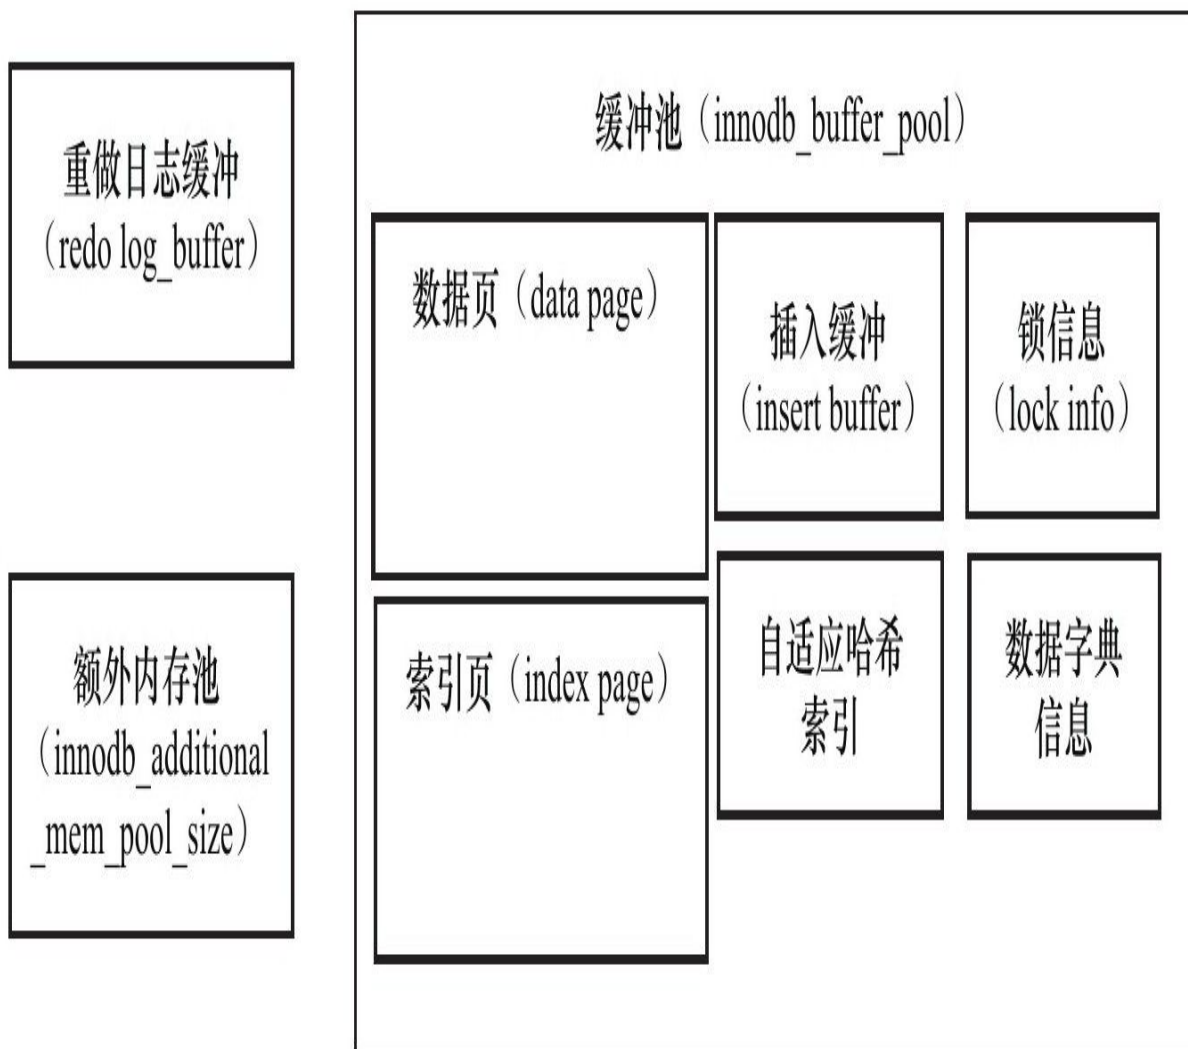
```
Variable_name:innodb_buffer_pool_size
```

```
Value:16106127360
```

```
1 row in set(0.00 sec)
```

InnoDB 数据库的默认配置参数如下所示。如果配置参数与默认值不同，则需要在配置文件中指定。如果配置参数与默认值相同，则不需要指定。如果配置参数与默认值不同，则需要使用“undo”来指定配置参数。如果配置参数与默认值不同，则需要使用“insert buffer”来指定配置参数。如果配置参数与默认值不同，则需要使用“adaptive hash index”来指定配置参数。如果配置参数与默认值不同，则需要使用“InnoDB”来指定配置参数。如果配置参数与默认值不同，则需要使用“lock info”来指定配置参数。

data dictionary
2-2 InnoDB



2-2 InnoDB

InnoDB 1.0.x
innodb_buffer_pool_instances

```
mysql> SHOW VARIABLES LIKE 'innodb_buffer_pool_instances'\G;
```

```
*****1. row*****
```

Variable_name:innodb_buffer_pool_instances

Value:1

1 row in set(0.00 sec)

innodb_buffer_pool_instances1
SHOW ENGINE INNODB STATUS

mysql>SHOW ENGINE INNODB STATUS\G;

*****1.row*****

Type:InnoDB

.....

INDIVIDUAL BUFFER POOL INFO

---BUFFER POOL 0

Buffer pool size 65535

Free buffers 65451

Database pages 84

Old database pages 0

Modified db pages 0

Pending reads 0

Pending writes:LRU 0,flush list 0 single page 0

Pages made young 0,not young 0

0.00 youngs/s,0.00 non-youngs/s

Pages read 84,created 0,written 1

9.33 reads/s,0.00 creates/s,0.11 writes/s

Buffer pool hit rate 764/1000,young-making rate 0/1000 not 0/1000

Pages read ahead 0.00/s,evicted without access 0.00/s,Random read ahead 0.00/s

LRU len:84,unzip_LRU len:0

I/O sum[0]:cur[0],unzip sum[0]:cur[0]

```
---BUFFER POOL 1

Buffer pool size 65536

Free buffers 65473

Database pages 63

Old database pages 0

Modified db pages 0

Pending reads 0

Pending writes:LRU 0,flush list 0 single page 0

Pages made young 0,not young 0

0.00 youngs/s,0.00 non-youngs/s

Pages read 63,created 0,written 0

7.00 reads/s,0.00 creates/s,0.00 writes/s

Buffer pool hit rate 500/1000,young-making rate 0/1000 not 0/1000

Pages read ahead 0.00/s,evicted without access 0.00/s,Random read ahead 0.00/s

LRU len:63,unzip_LRU len:0

I/O sum[0]:cur[0],unzip sum[0]:cur[0]
```

```
innodb_buffer_pool_instances2
SHOW ENGINE INNODB STATUS
---BUFFER POOL 0
```

```
MySQL 5.6
INNODB_BUFFER_POOL_STATS
```

```
mysql>SELECT P00L_ID,P00L_SIZE,
-FREE_BUFFERS,DATABASE_PAGES
-FROM INNODB_BUFFER_POOL_STATS\G;

*****1.row*****

P00L_ID:0

P00L_SIZE:65535
```

```

innodb_old_blocks_pct37
LRU37%3/8InnoDBmidpoint
oldnewnew

```

LRU LRU
LRU SQL
LRU
LRU InnoDB

InnoDB LRU
innodb_old_blocks_time mid LRU
SQL LRU

```
mysql>SET GLOBAL innodb_old_blocks_time=1000;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
#data or index scan operation
```

```
.....
```

```
mysql>SET GLOBAL innodb_old_blocks_time=0;
```

```
Query OK,0 rows affected(0.00 sec)
```

63% SQL

```
mysql>SET GLOBAL innodb_old_blocks_pct=20;
```

```
Query OK,0 rows affected(0.00 sec)
```

LRU LRU
Free Free
Free LRU LRU LRU
LRU old new page made
young innodb_old_blocks_time old new
page not made young SHOW ENGINE INNODB
STATUS LRU Free

```
mysql>SHOW ENGINE INNODB STATUS\G;
```

*****1.row*****

Type:InnoDB

Name:

Status:

=====

120725 22:04:25 INNODB MONITOR OUTPUT

=====

Per second averages calculated from the last 24 seconds

.....

Buffer pool size 327679

Free buffers 0

Database pages 307717

Old database pages 113570

Modified db pages 24673

Pending reads 0

Pending writes:LRU 0,flush list 0,single page 0

Pages made young 6448526,not young 0

48.75 youngs/s,0.00 non-youngs/s

Pages read 5354420,created 239625,written 3486063

55.68 reads/s,81.74 creates/s,955.88 writes/s

Buffer pool hit rate 1000/1000,young-making rate 0/1000 not 0/1000

.....

SHOW ENGINE INNODB STATUS
Buffer pool size 327 679
327679*16K 5GB
Free buffers
Free
Database pages
LRU
Free buffers
Database pages
Buffer pool size 2-2
Lock
Insert Buffer
LRU
LRU

pages made young LRU not young 0 youngs/s non-
young/s — Buffer pool hit
rate 100% 95% Buffer pool hit rate 95% LRU

SHOW ENGINE INNODB STATUS
InnoDB Per second averages
calculated from the last 24 seconds 24

InnoDB 1.2 INNODB_BUFFER_POOL_STATS

```
mysql> SELECT POOL_ID, HIT_RATE,  
- PAGES_MADE_YOUNG, PAGES_NOT_MADE_YOUNG  
- FROM information_schema.INNODB_BUFFER_POOL_STATS;  
  
*****1. row*****  
  
POOL_ID: 0  
  
HIT_RATE: 980  
  
PAGES_MADE_YOUNG: 450  
  
PAGES_NOT_MADE_YOUNG: 0
```

INNODB_BUFFER_PAGE_LRU LRU
LRU SPACE 1

```
mysql> SELECT TABLE_NAME, SPACE, PAGE_NUMBER, PAGE_TYPE  
- FROM INNODB_BUFFER_PAGE_LRU WHERE SPACE=1;  
  
+-----+-----+-----+-----+  
|TABLE_NAME|SPACE|PAGE_NUMBER|PAGE_TYPE|  
+-----+-----+-----+-----+  
  
|NULL|1|0|FILE_SPACE_HEADER|
```

|NULL|1|1|IBUF_BITMAP|

|NULL|1|2|INODE|

|test/t|1|3|INDEX|

+-----+-----+-----+-----+-----+

4 rows in set(0.00 sec)

InnoDB1.0.x16KB1KB2KB
4KB8KBLRU16KB
unzip_LRUSHOW ENGINE INNODB STATUS

mysqlSHOW ENGINE INNODB STATUS\G;

.....

Buffer pool hit rate 999/1000,young-making rate 0/1000 not 0/1000

Pages read ahead 0.00/s,evicted without access 0.00/s,Random read ahead 0.00/s

LRU len:1539,unzip_LRU len:156

I/O sum[0]:cur[0],unzip sum[0]:cur[0]

.....

LRU1539unzip_LRU156
LRUunzip_LRU

8KB2KB
unzip_LRU

unzip_LRU4KB

14KBunzip_LRU

2

38KBunzip_LRU

4. 4KB 2 4KB unzip_LRU

5. LRU 16KB 1 8KB 2 4KB unzip_LRU

information_schema INNODB_BUFFER_PAGE_LRU unzip_LRU

```
mysql> SELECT
```

```
-> TABLE_NAME, SPACE, PAGE_NUMBER, COMPRESSED_SIZE
```

```
-> FROM INNODB_BUFFER_PAGE_LRU
```

```
-> WHERE COMPRESSED_SIZE=0;
```

```
+-----+-----+-----+-----+
```

```
|TABLE_NAME|SPACE|PAGE_NUMBER|COMPRESSED_SIZE|
```

```
+-----+-----+-----+-----+
```

```
|sbtest/t|9|134|8192|
```

```
|sbtest/t|9|135|8192|
```

```
|sbtest/t|9|96|8192|
```

```
|sbtest/t|9|136|8192|
```

```
|sbtest/t|9|32|8192|
```

```
|sbtest/t|9|97|8192|
```

```
|sbtest/t|9|137|8192|
```

```
|sbtest/t|9|98|8192|
```

```
.....
```

LRU dirty page CHECKPOINT Flush LRU Flush

LRU Flush SHOW ENGINE INNODB STATUS
Modified db pages 24673

information_schema INNO_DB_BUFFER_PAGE_LRU
 LRU
 INNO_DB_BUFFER_PAGE_LRU
 OLDEST_MODIFICATION 0 SQL

```
mysql> SELECT TABLE_NAME,SPACE,PAGE_NUMBER,PAGE_TYPE
```

```
- FROM INNO_DB_BUFFER_PAGE_LRU
```

```
- WHERE OLDEST_MODIFICATION=0;
```

```
+-----+-----+-----+-----+
```

```
|TABLE_NAME|SPACE|PAGE_NUMBER|PAGE_TYPE|
```

```
+-----+-----+-----+-----+
```

```
|NULL|0|56|SYSTEM|
```

```
|NULL|0|0|FILE_SPACE_HEADER|
```

```
|test/t|1|3|INDEX|
```

```
|NULL|0|320|INODE|
```

```
|NULL|0|325|UNDO_LOG|
```

```
+-----+-----+-----+-----+
```

```
5 rows in set (0.00 sec)
```

5
 TABLE_NAME NULL

3.

2-2 InnoDB redo log
 buffer InnoDB
 innodb_log_buffer_size 8MB

```
mysql> SHOW VARIABLES LIKE 'innodb_log_buffer_size'\G;
```

```
*****1.row*****
```

Variable_name:innodb_log_buffer_size

Value:8388608

1 row in set(0.00 sec)

8MB 的缓冲池，用于存放redo log 的缓冲池，
innodb_log_buffer_size

Master Thread 主线程，用于管理缓冲池

缓冲池的初始化，包括缓冲池的分配和释放

缓冲池的大小为 1/2 的 innodb_buffer_pool_size

4. 缓冲池

DBA 管理缓冲池，包括缓冲池的分配和释放，
InnoDB 缓冲池的分配和释放，heap 缓冲池的分配和释放，
innodb_buffer_pool 缓冲池的分配和释放，frame buffer 缓冲池的分配和释放，
buffer control block 缓冲池的分配和释放，LRU 缓冲池的分配和释放，
InnoDB 缓冲池的分配和释放，
缓冲池

2.4 Checkpoint

データベースは、CPU、メモリ、ディスクなどのリソースを効率的に利用するために、定期的にデータをディスクに書き込み、メモリから削除する。この操作を「Checkpoint」と呼ぶ。この操作は、データベースのDML（Update、Delete）操作が行われるたびに実行される。

データベースは、Write Ahead Log（WAL）と呼ばれるログを保持する。このログは、データベースのACID（Atomicity, Consistency, Isolation, Durability）特性を確保するために使用される。

データベースは、定期的にデータをディスクに書き込み、メモリから削除する。この操作を「Checkpoint」と呼ぶ。この操作は、データベースのDML（Update、Delete）操作が行われるたびに実行される。

データベースは、定期的にデータをディスクに書き込み、メモリから削除する。

データベースは、定期的にデータをディスクに書き込み、メモリから削除する。

データベースは、定期的にデータをディスクに書き込み、メモリから削除する。この操作を「Checkpoint」と呼ぶ。この操作は、データベースのDML（Update、Delete）操作が行われるたびに実行される。3TBのMySQLデータベースは、3TBのデータを保持する。Oracle Exadataデータベースは、2TBのデータを保持する。

データベースは、定期的にデータをディスクに書き込み、メモリから削除する。この操作を「Checkpoint」と呼ぶ。この操作は、データベースのDML（Update、Delete）操作が行われるたびに実行される。DBA（Database Administrator）は、データベースの管理を行う。SA（System Administrator）は、データベースのシステム管理を行う。

データベースは、定期的にデータをディスクに書き込み、メモリから削除する。この操作を「Checkpoint」と呼ぶ。この操作は、データベースのDML（Update、Delete）操作が行われるたびに実行される。

データベースは、定期的にデータをディスクに書き込み、メモリから削除する。この操作を「Checkpoint」と呼ぶ。この操作は、データベースのDML（Update、Delete）操作が行われるたびに実行される。

データベースは、定期的にデータをディスクに書き込み、メモリから削除する。

❑ 数据库的Checkpoint

❑ 数据库的Checkpoint

数据库的Checkpoint是数据库在内存中保存的Checkpoint
数据库的Checkpoint是数据库在内存中保存的Checkpoint

数据库的Checkpoint是数据库在内存中保存的Checkpoint
Checkpoint是数据库在内存中保存的Checkpoint

数据库的Checkpoint是数据库在内存中保存的Checkpoint
数据库的Checkpoint是数据库在内存中保存的Checkpoint
数据库的Checkpoint是数据库在内存中保存的Checkpoint
数据库的Checkpoint是数据库在内存中保存的Checkpoint

InnoDB的LSN(Log Sequence Number)是数据库的LSN
LSN 8是数据库的LSN
LSN Checkpoint是数据库的LSN
SHOW ENGINE INNODB STATUS

```
mysql> SHOW ENGINE INNODB STATUS\G;
```

```
.....
```

```
---
```

```
LOG
```

```
---
```

```
Log sequence number 92561351052
```

```
Log flushed up to 92561351052
```

```
Last checkpoint at 92561351052
```

```
.....
```

InnoDB的Checkpoint是数据库的Checkpoint
Checkpoint是数据库的Checkpoint
Checkpoint是数据库的Checkpoint
Checkpoint是数据库的Checkpoint

☐Sharp Checkpoint

☐Fuzzy Checkpoint

Sharp Checkpoint
innodb_fast_shutdown=1

Sharp Checkpoint
InnoDB
Fuzzy Checkpoint

InnoDB
Fuzzy Checkpoint

☐Master Thread Checkpoint

☐FLUSH_LRU_LIST Checkpoint

☐Async/Sync Flush Checkpoint

☐Dirty Page too much Checkpoint

Master Thread
2.5
Master Thread
Checkpoint
InnoDB

FLUSH_LRU_LIST Checkpoint
InnoDB
LRU
100
InnoDB1.1.x
LRU
100
InnoDB
LRU
Checkpoint
LRU
FLUSH_LRU_LIST Checkpoint

MySQL 5.6
InnoDB1.2.x
Page
Cleaner
innodb_lru_scan_depth
LRU
1024

```
mysql> SHOW VARIABLES LIKE 'innodb_lru_scan_depth'\G;
```

```
*****1.row*****
```


InnoDB 1.2.x——MySQL 5.6Page Cleaner Thread

MySQLFlushLRUCheckpoint
Async/Sync FlushInnoDB
SHOW ENGINE INNODB STATUS

```
mysql>SHOW ENGINE INNODB STATUS\G;

*****1.row*****

Type:InnoDB

.....

LRU len:112902,unzip_LRU len:0

I/O sum[0]:cur[0],unzip sum[0]:cur[0]

Async Flush:0,Sync Flush:0,LRU List Flush:0,Flush List Flush:111736

.....

1 row in set (0.01 sec)
```

InnoDB9

CheckpointDirty Page too much
InnoDBCheckpoint
innodb_max_dirty_pages_pct

```
mysql>SHOW VARIABLES LIKE'innodb_max_dirty_pages_pct'\G;

*****1.row*****

Variable_name:innodb_max_dirty_pages_pct

Value:75

1 row in set (0.00 sec)
```

innodb_max_dirty_pages_pct7575%
CheckpointInnoDB 1.0.x

90□□□□□□75□

2.5 Master Thread

2.3 InnoDB Master Thread

2.5.1 InnoDB 1.0.x Master Thread

Master Thread loop loop
background loop flush loop suspend
loop Master Thread loop background loop
flush loop suspend loop

Loop — 10

```
void master_thread(){  
  
loop  
  
for(int i=0;i<10;i++){  
  
do thing once per second  
  
sleep 1 second if necessary  
  
}  
  
do things once per ten seconds  
  
goto loop;  
  
}
```

loop thread sleep 10
delay InnoDB

```
void master_thread(){
    goto loop;

loop
for(int i=0;i<10;i++){
    thread_sleep(1)//sleep 1 second

do log buffer flush to disk

    if(last_one_second_ios<5)

do merge at most 5 insert buffer

    if(buf_get_modified_ratio_pct<innodb_max_dirty_pages_pct)

do buffer pool flush 100 dirty page

    if(no user activity)

goto background loop
}

do things once per ten seconds
```

background loop:

do something

goto loop:

}

每隔10秒扫描一遍

每隔100秒扫描一遍

每隔5秒扫描一遍

每隔1秒扫描一遍

每隔1秒Undo扫描一遍

每隔100秒每隔10秒扫描一遍

每隔10秒InnoDB扫描一遍每隔10秒扫描IO扫描一遍200秒扫描一遍

InnoDB扫描一遍每隔10秒扫描IO扫描一遍100秒扫描一遍

InnoDB扫描一遍每隔10秒扫描一遍每隔10秒扫描一遍每隔10秒扫描一遍

每隔10秒扫描一遍InnoDB扫描一遍每隔10秒扫描一遍每隔10秒扫描一遍

每隔10秒扫描一遍

InnoDB扫描一遍full purge扫描一遍Undo扫描一遍

update delete扫描一遍consistent

read扫描一遍full purge扫描一遍InnoDB扫描一遍

扫描一遍undo扫描一遍

InnoDB扫描一遍InnoDB扫描一遍full purge

扫描一遍20undo

InnoDB扫描一遍buf_get_modified_ratio_pct

扫描一遍70%扫描一遍100扫描一遍70%扫描一遍10%

扫描一遍

main loop扫描一遍

```

void master_thread(){

goto loop;

loop

for(int i=0;i<10;i++){

thread_sleep(1)//sleep 1 second

do log buffer flush to disk

if(last_one_second_ios<5)

do merge at most 5 insert buffer

if(buf_get_modified_ratio_pct<innodb_max_dirty_pages_pct)

do buffer pool flush 100 dirty page

if(no user activity)

goto backgroud loop

}

if(last_ten_second_ios<200)

do buffer pool flush 100 dirty page

do merge at most 5 insert buffer

do log buffer flush to disk

do full purge

if(buf_get_modified_ratio_pct<70%)

do buffer pool flush 100 dirty page

else

buffer pool flush 10 dirty page

goto loop

background loop:

do something

goto loop:

}

```

background loop
shutdown
background loop

□□□□□Undo□□□□□□

□□□20□□□□□□□□□□

□□□□□□□□□□□□

□□□□□100□□□□□□□□□□□□flush loop□□□□□

□flush loop□□□□□□□□□□□InnoDB□□□□□□□suspend__loop□□
Master Thread□□□□□□□□□□□□enable□□InnoDB□□□□□□□□
□□□□InnoDB□□□□□□□□□Master Thread□□□□□□□□□□

□□□Master Thread□□□□□□□□□□

```
void master_thread(){  
  
    goto loop;  
  
loop:  
  
    for(int i=0;i<10;i++){  
  
        thread_sleep(1)//sleep 1 second  
  
        do log buffer flush to disk  
  
        if(last_one_second_ios<5)  
  
            do merge at most 5 insert buffer  
  
            if(buf_get_modified_ratio_pct<innodb_max_dirty_pages_pct)  
  
                do buffer pool flush 100 dirty page  
  
                if(no user activity)  
  
                    goto backgroud loop  
  
            }  
  
            if(last_ten_second_ios<200)  
  
                do buffer pool flush 100 dirty page  
  
                do merge at most 5 insert buffer  
  
                do log buffer flush to disk  
  
                do full purge
```

```
if(buf_get_modified_ratio_pct<70%)

do buffer pool flush 100 dirty page

else

buffer pool flush 10 dirty page

goto loop

background loop:

do full purge

do merge 20 insert buffer

if not idle:

goto loop:

else:

goto flush loop

flush loop:

do buffer pool flush 100 dirty page

if(buf_get_modified_ratio_pct<innodb_max_dirty_pages_pct)

goto flush loop

goto suspend loop

suspend loop:

suspend_thread()

waiting event

goto loop;

}
```

innodb_max_dirty_pages_pct 75 Google 80
IO

InnoDB 1.0.x innodb_adaptive_flushing
innodb_max_dirty_pages_pct
innodb_max_dirty_pages_pct 100
innodb_adaptive_flushing InnoDB
buf_flush_get_desired_flush_rate
buf_flush_get_desired_flush_rate redo
log
innodb_max_dirty_pages_pct

full purge 20 Undo InnoDB
1.0.x innodb_purge_batch_size full
purge Undo 20

```
mysql>SHOW VARIABLES LIKE'innodb_purge_batch_size'\G;

*****1.r0w*****

Variable_name:innodb_purge_batch_size

Value:20

mysql>SET GLOBAL innodb_purge_batch_size=50;

Query OK,0 rows affected(0.00 sec)
```

InnoDB 1.0.x Master Thread

```
void master_thread(){

goto loop;

loop

for(int i=0;i<10;i++){

thread_sleep(1)//sleep 1 second

do log buffer flush to disk
```

```

if(last_one_second_ios[]5%innodb_io_capacity)

do merge 5%innodb_io_capacity insert buffer

if(buf_get_modified_ratio_pct[]innodb_max_dirty_pages_pct)

do buffer pool flush 100%innodb_io_capacity dirty page

else if enable adaptive flush

do buffer pool flush desired amount dirty page

if(no user activity)

goto backgroud loop

}

if(last_ten_second_ios[]innodb_io_capacity)

do buffer pool flush 100%innodb_io_capacity dirty page

do merge 5%innodb_io_capacity insert buffer

do log buffer flush to disk

do full purge

if(buf_get_modified_ratio_pct[]70%)

do buffer pool flush 100%innodb_io_capacity dirty page

else

do buffer pool flush 10%innodb_io_capacity dirty page

goto loop

background loop:

do full purge

do merge 100%innodb_io_capacity insert buffer

if not idle:

goto loop:

else:

goto flush loop

flush loop:

do buffer pool flush 100%innodb_io_capacity dirty page

if(buf_get_modified_ratio_pct[]innodb_max_dirty_pages_pct)

go to flush loop

```

```

goto suspend loop

suspend loop:

suspend_thread()

waiting event

goto loop;

}

```

InnoDB 1.0.x Master Thread
 InnoDB Master Thread

InnoDB 1.0.x SHOW ENGINE INNODB STATUS
 Master Thread

```
mysql> SHOW ENGINE INNODB STATUS\G;
```

```
*****1.row*****
```

```
Type:InnoDB
```

```
Name:
```

```
Status:
```

```
=====
```

```
090921 14:24:56 INNODB MONITOR OUTPUT
```

```
=====
```

```
Per second averages calculated from the last 6 seconds
```

```
-----
```

```
BACKGROUND THREAD
```

```
-----
```

```
srv_master_thread loops:45 1_second,45 sleeps,4 10_second,6 background,6 flush
```

```
srv_master_thread log flush and writes:45 log writes only:69
```

```
.....
```

45sleep4510
41:10background loop6flush loop6
MySQL

```
mysql>show engine innodb status\G;
```

```
*****1. row*****
```

```
Type:InnoDB
```

```
Name:
```

```
Status:
```

```
=====
```

```
091009 10:14:34 INNODB MONITOR OUTPUT
```

```
=====
```

```
Per second averages calculated from the last 42 seconds
```

```
-----
```

```
BACKGROUND THREAD
```

```
-----
```

```
srv_master_thread loops:2188 1_second,1537 sleeps,218 10_second,2 background,2 flush
```

```
srv_master_thread log flush and writes:1777 log writes only:5816
```

```
.....
```

2188sleep1537
InnoDB11_second
sleeps

[1] <http://code.google.com/p/google-mysql-tools/wiki/InnoblockItpDisk>

2.5.3 InnoDB 1.2.x의 Master Thread

InnoDB 1.2.x의 Master Thread는 InnoDB 1.2.x의 Master Thread가 InnoDB 1.2.x의 Master Thread를 실행하는 동안

```
if InnoDB is idle
    srv_master_do_idle_tasks();
else
    srv_master_do_active_tasks();
```

srv_master_do_idle_tasks()는 10초마다
srv_master_do_active_tasks()는 InnoDB의 Master Thread가 Page Cleaner Thread와 Master Thread가 InnoDB의 Master Thread를 실행하는 동안

2.6 InnoDB引擎

InnoDB引擎特性

❑ 支持Insert Buffer

❑ 支持Double Write

❑ 支持Adaptive Hash Index

❑ 支持IO Async IO

❑ 支持Flush Neighbor Page

支持InnoDB引擎特性

2.6.1 引擎特性

1.Insert Buffer

Insert Buffer是InnoDB引擎中用于处理批量插入数据的一种机制。当插入的数据量较大时，InnoDB引擎会将数据暂时存储在Insert Buffer中，然后批量写入磁盘。这样可以减少磁盘I/O次数，提高插入效率。

InnoDB引擎支持Primary Key，即主键。主键是表中唯一标识每一行数据的字段。在SQL中，主键通常用PRIMARY KEY来定义。

```
CREATE TABLE t(  
  a INT AUTO_INCREMENT,  
  b VARCHAR(30),  
  PRIMARY KEY(a)  
);
```

如果a是主键，a不能为NULL，且a是AUTO_INCREMENT，那么a的值是唯一的，且a的值是递增的。

如果a是主键，a不能为NULL，且a是UUID，那么a的值是唯一的，且a的值是随机的。

如果a是主键，a不能为NULL，且a是secondary index，那么a的值是唯一的，且a的值是随机的。SQL中，主键的索引是唯一的，且是自动的。

```
CREATE TABLE t(  
  a INT AUTO_INCREMENT,  
  b VARCHAR(30),  
  PRIMARY KEY(a),  
  key(b)  
);
```

如果a是主键，a不能为NULL，且a是B+，那么a的值是唯一的，且a的值是随机的。B+树是一种平衡的树结构，用于存储大量的数据。

如果a是主键，a不能为NULL，且a是“”索引，那么a的值是唯一的，且a的值是随机的。索引是一种数据结构，用于快速查找数据。

InnoDB的Insert Buffer是一种缓冲机制，用于处理大量的插入操作。Insert Buffer是一种缓冲机制，用于处理大量的插入操作。Insert Buffer是一种缓冲机制，用于处理大量的插入操作。merge是一种合并操作，用于将多个数据合并成一个。

Insert Buffer是一种缓冲机制，用于处理大量的插入操作。

□□□□□□□unique□□□

[illegible]

SHOW ENGINE INNODB STATUS

```
*****1.row*****
```

Name:

Status:

=====

```
100727 22:21:48 INNODB MONITOR OUTPUT
```

Per second averages calculated from the last 44 seconds

●●●●●●

.....

INSERT BUFFER AND ADAPTIVE HASH INDEX

```
Ibuf:size 7545,free list len 3790,seg size 11336,
```

```
8075308 inserts,7540969 merged recs,2246304 merges
```

.....

END OF INNODB MONITOR OUTPUT

=====

1 row in set(0.00 sec)

seg size Insert Buffer 11336×16KB 177MB free
list len size 2
Inserts merged recs
merges merges:merged recs
1:3 IO 2/3

Insert Buffer
innodb_buffer_pool 1/2 InnoDB
insert buffer

```
/**Buffer pool size per the maximum insert buffer size*/
```

```
#define IBUF_POOL_SIZE_PER_MAX_SIZE 2
```

```
ibuf->max_size=buf_pool_get_curr_size()/UNIV_PAGE_SIZE
```

```
/IBUF_POOL_SIZE_PER_MAX_SIZE;
```

Percona patch
Percona
IBUF_POOL_SIZE_PER_MAX_SIZE
IBUF_POOL_SIZE_PER_MAX_SIZE 3 1/3

2.Change Buffer

InnoDB 1.0.x Change Buffer Insert Buffer
InnoDB DML — INSERT DELETE UPDATE
Insert Buffer Delete Buffer Purge buffer

Insert Buffer Change Buffer

UPDATE

Delete Buffer UPDATE Purge Buffer
UPDATE InnoDB
innodb_change_buffering Buffer
inserts deletes purges changes all none inserts deletes
purges changes inserts deletes all
none all

InnoDB 1.2.x innodb_change_buffer_max_size
Change Buffer

```
mysql>SHOW VARIABLES LIKE'innodb_change_buffer_max_size'\G;
```

```
*****1.row*****
```

```
Variable_name:innodb_change_buffer_max_size
```

```
Value:25
```

```
1 row in set(0.00 sec)
```

innodb_change_buffer_max_size 25 1/4
50

MySQL 5.5 SHOW ENGINE INNODB STATUS

```
mysql>SHOW ENGINE INNODB STATUS\G;
```

```
*****1.row*****
```

```
Type:InnoDB
```

```
.....
```

```
-----  
INSERT BUFFER AND ADAPTIVE HASH INDEX
```

```
-----  
Ibuf:size 1,free list len 34397,seg size 34399,10875 merges
```

```
merged operations:
```

```
insert 20462,delete mark 20158,delete 4215
```

discarded operations:

insert 0,delete mark 0,delete 0

.....

merged operations discarded operation
Change Buffer insert Insert Buffer delete mark
Delete Buffer delete Purge Buffer discarded operations
Change Buffer merge merge
merge

3.Insert Buffer

Insert Buffer
Insert Buffer

Insert Buffer B+ MySQL 4.1
Insert Buffer B+ Insert Buffer
B+ Insert Buffer B+
ibdata1 ibd CHECK TABLE
Insert Buffer ibd
REPAIR TABLE

Insert Buffer B+ search
key 2-3

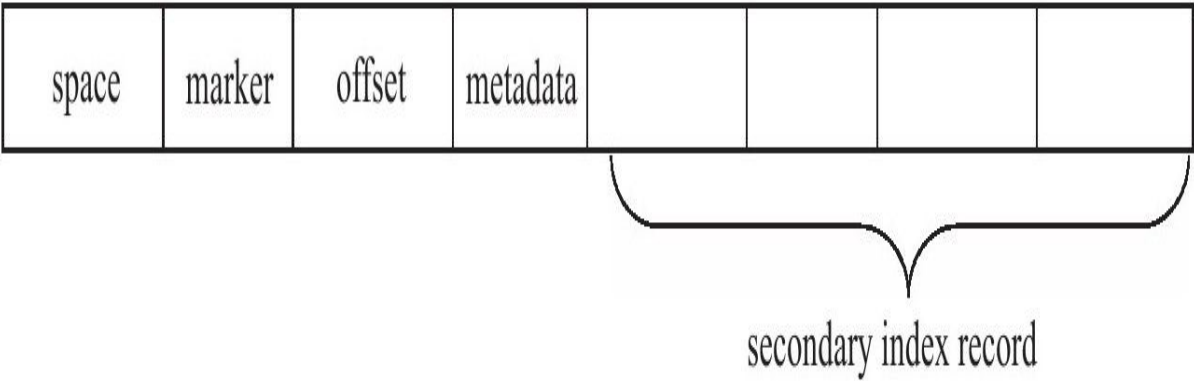
space	marker	offset
-------	--------	--------

2-3 Insert Buffer search key

search key 9 space id InnoDB
space id space id space 4
marker 1 Insert Buffer offset
4

space offset InnoDB
search key Insert Buffer B+
Insert Buffer B+

Insert Buffer B+ 2-4



2-4 Insert Buffer

space marker page_no 9 4
metadata 4 2-2

表 2-2 metadata 字段存储的内容

名 称	字 节
IBUF_REC_OFFSET_COUNT	2
IBUF_REC_OFFSET_TYPE	1
IBUF_REC_OFFSET_FLAGS	1

IBUF_REC_OFFSET_COUNT 记录 Insert Buffer 在 InnoDB 1.0.x 版本 Change Buffer 中 replay 的总次数

记录 Insert Buffer 在 5 个 redo log 中 replay 的总次数
Insert Buffer B+ 树中记录 13 个 redo log

记录 Insert Buffer 在 space 中 page_no 的记录
Insert Buffer B+ 树中 Merge Insert Buffer 的记录
记录 Insert Buffer 在 space 中 page_no 的记录
Insert Buffer Bitmap

记录 Insert Buffer Bitmap 中 16384 个 bit 的记录
Extent 记录 Insert Buffer Bitmap 中 16384 个 bit 的记录
Insert Buffer Bitmap

记录 Insert Buffer Bitmap 中 4 个 bit 的记录 2-3 个 redo log

表 2-3 每个辅助索引页在 Insert Buffer Bitmap 中存储的信息

名 称	大小 (bit)	说 明
IBUF_BITMAP_FREE	2	表示该辅助索引页中的可用空间数量，可取值为： <input type="checkbox"/> 0 表示无可用剩余空间 <input type="checkbox"/> 1 表示剩余空间大于 1/32 页（512 字节） <input type="checkbox"/> 2 表示剩余空间大于 1/16 页 <input type="checkbox"/> 3 表示剩余空间大于 1/8 页
IBUF_BITMAP_BUFFERED	1	1 表示该辅助索引页有记录被缓存在 Insert Buffer B+ 树中
IBUF_BITMAP_IBUF	1	1 表示该页为 Insert Buffer B+ 树的索引页

4.Merge Insert Buffer

- ☐ Insert/Change Buffer ☐ B+ ☐
- ☐ B+ ☐ Insert Buffer ☐
- ☐ merge ☐
- ☐ Merge Insert Buffer ☐
- ☐
- ☐ Insert Buffer Bitmap ☐
- ☐ Master Thread ☐

SELECT
Insert Buffer Bitmap
Insert Buffer B+
Insert Buffer B+

Insert Buffer Bitmap
1/32
1/32
Insert Buffer B+

Master Thread
Master Thread
10
Merge Insert Buffer
merge

Master Thread
merge
srv_innodb_io_capacity
InnoDB

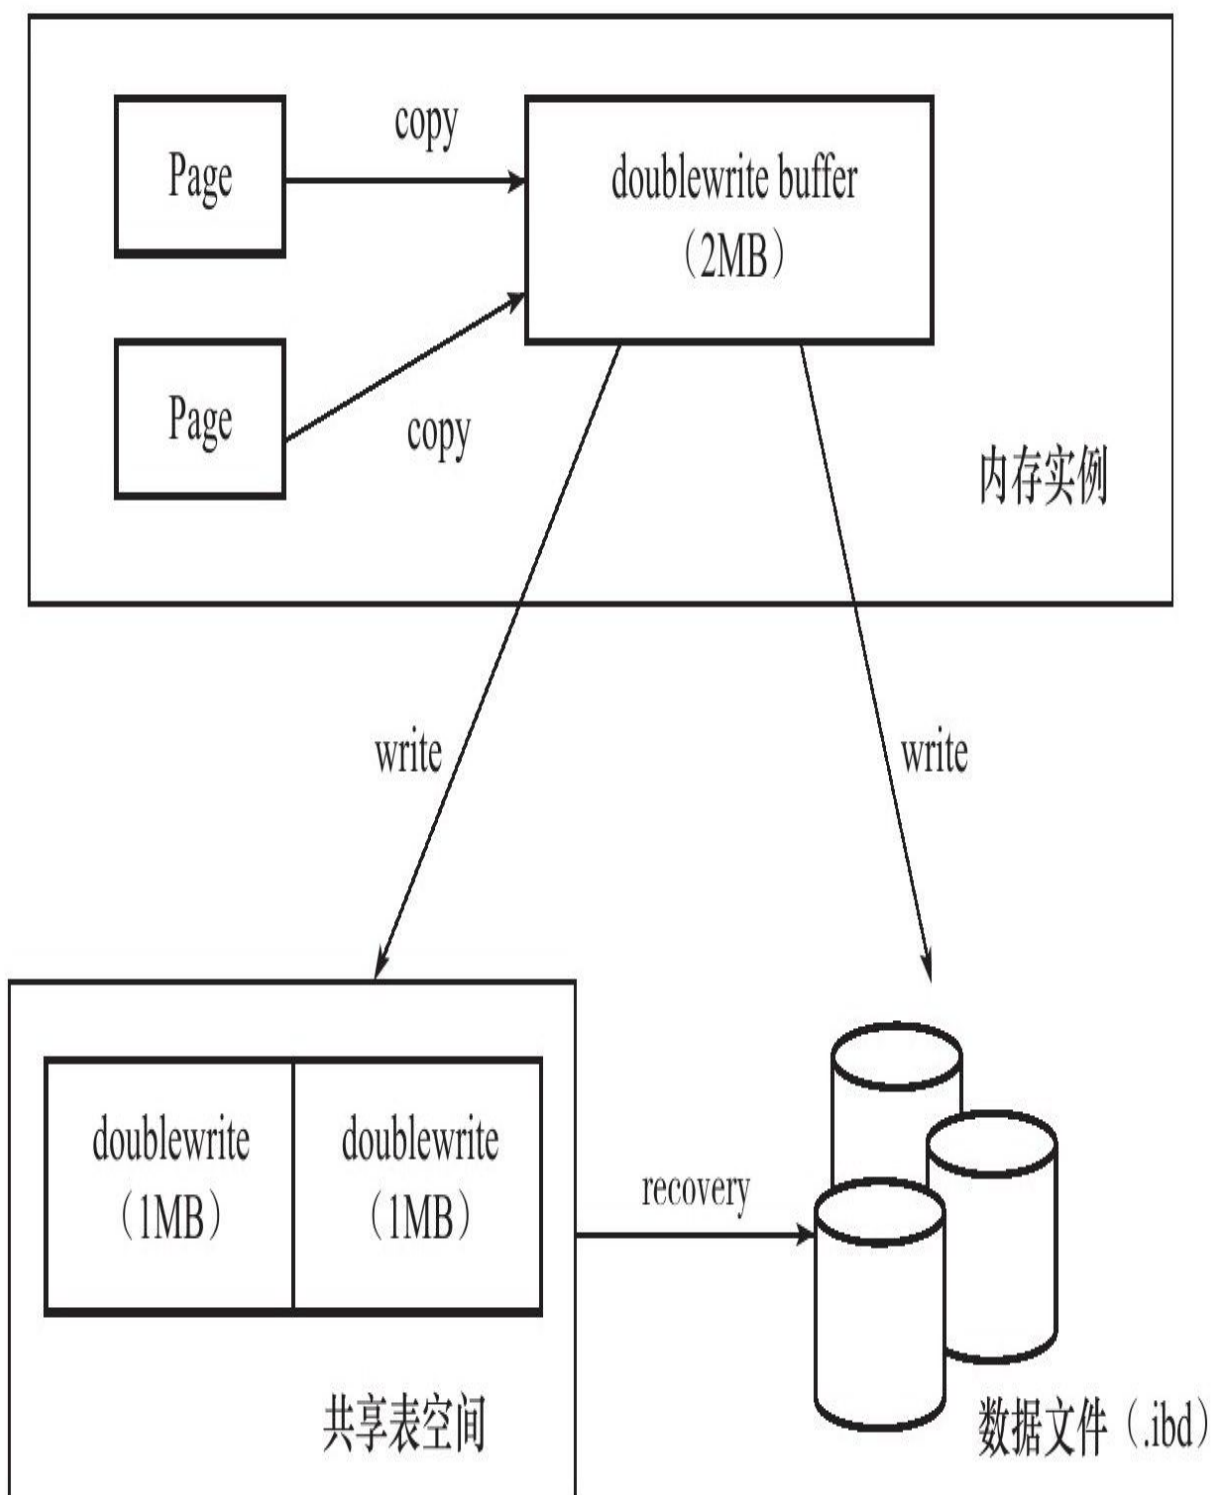
Insert Buffer B+
space
offset
Insert Buffer
InnoDB
Insert Buffer B+
space
merge
merge
Insert/Change Buffer

2.6.2 问题

Insert Buffer InnoDB doublewrite InnoDB

InnoDB 16KB 4KB partial page write InnoDB doublewrite

DBA InnoDB 800 'aaaa' apply doublewrite InnoDB doublewrite 2-5



□ 2-5 InnoDB□□□doublewrite□□

doublewrite buffer 2MB
128 extent 2MB
memcpy doublewrite
buffer doublewrite buffer 1MB
fsync doublewrite
doublewrite
doublewrite buffer
doublewrite

```
mysql>SHOW GLOBAL STATUS LIKE'innodb_dblwr'\G;
```

```
*****1.row*****
```

```
Variable_name:Innodb_dblwr_pages_written
```

```
Value:6325194
```

```
*****2.row*****
```

```
Variable_name:Innodb_dblwr_writes
```

```
Value:100399
```

```
2 rows in set(0.00 sec)
```

doublewrite 6 325 194 100 399
64:1
Innodb_dblwr_pages_written:Innodb_dblwr_writes 64:1

InnoDB
doublewrite
doublewrite

```
090924 11:36:32 mysqld restarted
```

```
090924 11:36:33 InnoDB:Database was not shut down normally!
```

```
InnoDB:Starting crash recovery.
```

```
InnoDB:Reading tablespace information from the.ibd files...
```

```
InnoDB:Crash recovery may have failed for some.ibd files!
```

InnoDB:Restoring possible half-written data pages from the doublewrite

InnoDB:buffer...

MySQLSHOW GLOBAL STATUS
Innodb_buffer_pool_pages_flushed
doublewrite
Innodb_dblwr_pages_writtenMySQL 5.5.24
Innodb_buffer_pool_pages_flushed
Innodb_dblwr_pages_written2BugMySQL5.5.24
Innodb_dblwr_pages_writtenMySQL

skip_innodb_doublewrite
slave server
slaves erverRAID0
master serverdoublewrite

ZFS
doublewrite

2.6.3 B+ Tree

hash 索引的查找复杂度为 $O(1)$ ，而 B+ 索引的查找复杂度为 $O(\log_3 4)$ 。

InnoDB 使用 Adaptive Hash Index (AHI) 索引。AHI 索引是 B+ 索引的变种，它使用哈希表来存储索引键。InnoDB 使用 AHI 索引来加速查询。

AHI 索引的查找复杂度为 $O(1)$ ，而 B+ 索引的查找复杂度为 $O(\log_3 4)$ 。

WHERE a=xxx

WHERE a=xxx and b=xxx

InnoDB 使用 AHI 索引来加速查询。

100

$N = N * 1/16$

InnoDB 使用 AHI 索引来加速查询。self-tuning DBA

SHOW ENGINE INNODB STATUS AHI

```
mysql> SHOW ENGINE INNODB STATUS\G;
```

```
*****1.row*****
```

```
Status:
```

```
=====
```

090922 11:52:51 INNODB MONITOR OUTPUT

=====

Per second averages calculated from the last 15 seconds

.....

INSERT BUFFER AND ADAPTIVE HASH INDEX

Ibuf:size 2249,free list len 3346,seg size 5596,

374650 inserts,51897 merged recs,14300 merges

Hash table size 4980499,node heap has 1246 buffer(s)

1640.60 hash searches/s,3709.46 non-hash searches/s

.....

#####AHI#####AHI#####AHI#####
#####SELECT*FROM table WHERE
index_col='xxx'#####
non-hash searches/s#####
hash searches:non-hash
searches#####

##AHI##InnoDB#####SHOW
ENGINE INNODB STATUS#####innodb_adaptive_hash_index
#####AHI#####

2.6.4 AIO

MySQL 5.6.16 版本开始支持异步 IO (Asynchronous IO, AIO) 功能，
InnoDB 支持 AIO。

AIO 与 Sync IO 不同，Sync IO 是同步的，而 AIO 是异步的。
SQL 语句执行时，InnoDB 会调用 AIO 接口来读写数据。
AIO 接口在 MySQL 5.6.16 版本中引入，位于 `libaio` 库中。
AIO 接口在 MySQL 5.6.16 版本中引入，位于 `libaio` 库中。

AIO 接口在 MySQL 5.6.16 版本中引入，位于 `libaio` 库中。
AIO 接口在 MySQL 5.6.16 版本中引入，位于 `libaio` 库中。

8 6 8 7 8 8

16KB 的 IO 大小，3 个 IO 大小，AIO 接口在 MySQL 5.6.16
版本中引入，位于 `libaio` 库中。AIO 接口在 MySQL 5.6.16
版本中引入，位于 `libaio` 库中。

Linux 系统使用 `iostat` 命令查看 IO 统计信息。

```
avg-cpu:  %user%nice%system%iowait%steal%idle
```

```
4.70  0.00  1.60 13.20  0.00 80.50
```

```
Device:rrqm/s wrqm/s r/s w/s rMB/s wMB/s avgrq-sz avgqu-sz await  svctm%util
```

```
sdc 3905.67 172.00 6910.33 466.67 168.81 18.15 51.91 19.17 2.59 0.13 97.73
```

InnoDB 1.1.x 版本支持 AIO 功能，InnoDB 1.1.x 版本支持 AIO 功能。
InnoDB 1.1.x 版本支持 AIO 功能，InnoDB 1.1.x 版本支持 AIO 功能。
InnoDB 1.1.x 版本支持 AIO 功能，InnoDB 1.1.x 版本支持 AIO 功能。

```
/usr/local/mysql/bin/mysqld:error while loading shared libraries:libaio.so.1:cannot open shared object file:No such file or directory
```

MySQL Native AIO Windows Linux
Native AIO Mac OSX
MySQL

innodb_use_native_aio Native AIO Linux
ON

```
mysql>SHOW VARIABLES LIKE'innodb_use_native_aio'\G;

*****1.row*****

Variable_name:innodb_use_native_aio

Value:ON

1 row in set(0.00 sec)
```

MySQL Native AIO InnoDB
Native AIO 75%

InnoDB read ahead AIO
AIO

2.6.5 性能调优

InnoDB 的 Flush Neighbor Page 操作会遍历整个 InnoDB 的 extent，这会带来大量的 IO 操作，如果 InnoDB 的 extent 非常大，那么 Flush Neighbor Page 操作就会带来大量的 IO 操作，这会严重影响数据库的性能。

❑ 可以通过设置 `innodb_flush_neighbors` 参数来优化性能。

❑ `innodb_flush_neighbors` 参数的默认值是 1，表示每次 flush 操作都会 flush 整个 extent。

InnoDB 1.2.x 版本中，`innodb_flush_neighbors` 参数的默认值是 1，表示每次 flush 操作都会 flush 整个 extent。如果将 `innodb_flush_neighbors` 参数设置为 0，那么每次 flush 操作只会 flush 当前页，这可以大大减少 IO 操作，从而提高数据库的性能。

2.7 数据库崩溃

InnoDB MySQL 数据库崩溃 InnoDB 数据库崩溃 MySQL 数据库崩溃 InnoDB 数据库崩溃

innodb_fast_shutdown 数据库崩溃 InnoDB 数据库崩溃 0 1 2 数据库崩溃 1

0 MySQL 数据库崩溃 InnoDB 数据库崩溃 full purge merge insert buffer 数据库崩溃 InnoDB 数据库崩溃 0 数据库崩溃

1 innodb_fast_shutdown 数据库崩溃 full purge merge insert buffer 数据库崩溃

2 full purge merge insert buffer 数据库崩溃 MySQL 数据库崩溃 recovery

MySQL 数据库崩溃 “kill” 数据库崩溃 MySQL 数据库崩溃 innodb_fast_shutdown 2 MySQL 数据库崩溃 InnoDB 数据库崩溃

innodb_force_recovery 数据库崩溃 InnoDB 数据库崩溃 0 数据库崩溃 corruption MySQL 数据库崩溃 crash 数据库崩溃

alter table 数据库崩溃 InnoDB 数据库崩溃 数据库崩溃

innodb_force_recovery 6 1 6 数据库崩溃

❑1(SRV_FORCE_IGNORE_CORRUPT)强制忽略corrupt

❑2(SRV_FORCE_NO_BACKGROUND)强制Master Thread不运行
❑Master Thread不运行full purge强制crash

❑3(SRV_FORCE_NO_TRX_UNDO)强制不运行Undo

❑4(SRV_FORCE_NO_IBUF_MERGE)强制不运行InnoDB的缓冲池合并

❑5(SRV_FORCE_NO_UNDO_LOG_SCAN)强制不运行Undo Log
❑InnoDB不运行Undo Log扫描

❑6(SRV_FORCE_NO_LOG_REDO)强制不运行redo log

强制InnoDB的innodb_force_recovery=0强制InnoDB不运行select
create drop insert update delete DML

强制InnoDB的innodb_force_recovery=1 000强制
InnoDB不运行redo log

```
mysql>START TRANSACTION;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>UPDATE Profile SET password='';
```

```
Query OK,9587770 rows affected(7 min 55.73 sec)
```

```
Rows matched:9999248 Changed:9587770 Warnings:0
```

START TRANSACTION强制不运行auto commit
❑UPDATE强制不运行UNDO
❑undo log强制不运行kill
MySQL强制不运行

```
[root@nineyou0-43]#ps -ef|grep mysqld
```

```
root 28007 1 0 13:40 pts/1 00:00:00/bin/sh./bin/mysqld_safe--datadir=/usr/local/mysql/data--pid-  
file=/usr/local/mysql/data/nineyou0-43.pid
```

```
mysql 28045 28007 42 13:40 pts/1 00:04:23/usr/local/mysql/bin/mysqld--basedir=/usr/local/mysql--datadir=/usr/local/mysql/data-  
-user=mysql--pid-file=/usr/local/mysql/data/nineyou0-43.pid--skip-external-locking--port=3306--socket=/tmp/mysql.sock
```

root 28110 26963 0 13:50 pts/11 00:00:00 grep mysqld

[root@nineyou0-43~]#kill-9 28007

[root@nineyou0-43~]#kill-9 28045

killMySQLUPDATE
err

090922 13:40:20 InnoDB:Started;log sequence number 6 2530474615

InnoDB:Starting in background the rollback of uncommitted transactions

090922 13:40:20 InnoDB:Rolling back trx with id 0 5281035,8867280 rows to undo

InnoDB:Progress in percents:1090922 13:40:20

090922 13:40:20[Note]/usr/local/mysql/bin/mysqld:ready for connections.

Version:'5.0.45-log'socket:'/tmp/mysql.sock'port:3306 MySQL Community Server(GPL)

2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46
47 48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80 81 82 83 84 85 86 87 88
89 90 91 92 93 94 95 96 97 98 99 100

InnoDB:Rolling back of trx id 0 5281035 completed

090922 13:49:21 InnoDB:Rollback of non-prepared transactions completed

innodb_force_recovery0InnoDB
8867280
9

MySQLinnodb_force_recovery
3InnoDB

090922 14:26:23 InnoDB:Started;log sequence number 7 2253251193

InnoDB:!!!innodb_force_recovery is set to 3!!!

090922 14:26:23[Note]/usr/local/mysql/bin/mysqld:ready for connections.

Version:'5.0.45-log'socket:'/tmp/mysql.sock'port:3306 MySQL Community Server(GPL)

~~~~~“!!!”~InnoDB~~~~innodb\_force\_recovery~~~3~~~~~  
~~~~~  
~~

2.8 〇〇

InnoDB InnoDB InnoDB
 InnoDB InnoDB InnoDB
 InnoDB “ ” MySQL
 InnoDB

MySQL InnoDB 3 MySQL MySQL InnoDB InnoDB

3 問

MySQLとInnoDBについて、正しいものを2つ選べ。

☐ MySQLは、デフォルトでInnoDBエンジンを使用する。

☐ MySQLは、デフォルトでMyISAMエンジンを使用する。

☐ socketファイルは、UNIX系OSでしか使用されない。

☐ pidファイルは、MySQLの起動時に生成される。

☐ MySQLは、デフォルトで32ビットの整数を使用する。

☐ MySQLは、デフォルトでInnoDBエンジンを使用する。

3.1 □□□□

```
1 MySQL
MySQL
mysql--help|grep my.cnf
```

MySQLOracleOracle
mountMySQLMySQL
MySQL
MySQLmysql

```
090922 16:25:52 mysqld started
```

```
090922 16:25:53 InnoDB:Started;log sequence number 8 2801063211
```

```
InnoDB:!!!innodb force recovery is set to 1!!!
```

```
090922 16:25:53[ERROR]Fatal error:Can't open and lock privilege tables:Table'mysql.host'doesn't exist
```

```
090922 16:25:53 mysqld ended
```

MySQLmysqlMySQL

MySQL
emacs

3.1.1

```

mysql> SET GLOBAL innodb_buffer_pool_size=1G;
mysql> SHOW VARIABLES LIKE 'innodb_buffer_pool_size';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| innodb_buffer_pool_size | 1G    |
+-----+-----+

```

information_schema GLOBAL_VARIABLES

```
mysql>SELECT*FROM
```

```
- GLOBAL_VARIABLES
```

```
- WHERE VARIABLE_NAME LIKE'innodb_buffer%\G;
```

```
*****1.row*****
```

```
VARIABLE_NAME:INNODB_BUFFER_POOL_SIZE
```

```
VARIABLE_VALUE:1073741824
```

```
1 row in set(0.00 sec)
```

```
mysql>SHOW VARIABLES LIKE'innodb_buffer%\G;
```

```
*****1.row*****
```

```
Variable_name:innodb_buffer_pool_size
```

```
Value:1073741824
```

```
1 row in set(0.00 sec)
```

GLOBAL_VARIABLES
SHOW VARIABLES MySQL

Oracle undocumented parameter

Oracle“”SQL ServerDBAMySQL
OracleSQL Server
MySQL

3.1.2 变量

MySQL 变量

动态变量

静态变量

MySQL 变量分为静态变量和动态变量。静态变量只能在 MySQL 启动时通过 `SET` 命令设置，动态变量可以在 MySQL 运行过程中通过 `SET` 命令设置。

SET

`|[global|session]system_var_name=expr`

`|[@@global.|@@session.|@@]system_var_name=expr`

MySQL 变量分为静态变量和动态变量。静态变量只能在 MySQL 启动时通过 `SET` 命令设置，动态变量可以在 MySQL 运行过程中通过 `SET` 命令设置。静态变量包括 `autocommit`、`binlog_cache_size`、`read_buffer_size` 等。

```
mysql> SET read_buffer_size=524288;
```

```
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> SELECT @@session.read_buffer_size\G;
```

```
*****1.row*****
```

```
@@session.read_buffer_size:524288
```

```
1 row in set (0.00 sec)
```

```
mysql> SELECT @@global.read_buffer_size\G;
```

```
*****1.row*****
```

```
@@global.read_buffer_size:2093056
```

```
1 row in set (0.00 sec)
```

read_buffer_size 2MB 512KB
read_buffer_size 2MB MySQL
read_buffer_size 2MB 512KB set global|session
SET @@global|@@session

```
mysql> SET @@global.read_buffer_size=1048576;
```

```
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> SELECT @@session.read_buffer_size\G;
```

```
*****1.row*****
```

```
@@session.read_buffer_size:524288
```

```
1 row in set (0.00 sec)
```

```
mysql> SELECT @@global.read_buffer_size\G;
```

```
*****1.row*****
```

```
@@global.read_buffer_size:1048576
```

```
1 row in set (0.00 sec)
```

read_buffer_size 1MB read_buffer_size
512KB
MySQL MySQL
MySQL
MySQL Dynamic System
Variables

```
mysql> SET GLOBAL datadir='/db/mysql';
```

```
ERROR 1238 (HY000): Variable 'datadir' is a read only variable
```

3.2 〇〇〇〇

MySQL MySQL

error log

binlog

slow query log

 $\square\square\square\square\square\log\square$

DBA MySQL

3.2.1 四則運算

```

MySQL
MySQL DBA
SHOW VARIABLES LIKE 'log_error'

```

```
mysql> SHOW VARIABLES LIKE 'log_error'\G;
```

```
*****1.row*****
```

Variable_name:log_error

Value:/mysql_data_2/stargazer.log

```
1 row in set (0.00 sec)
```

```
mysql>system hostname
```

stargazer

```

stargazer
startgazer.err
MySQL

```

```
[root@nineyou0-43 data]#tail-n 50 nineyou0-43.err
```

```
090924 11:31:18 mysqld started
```

```
090924 11:31:18 InnoDB:Started;log sequence number 8 2801063331
```

```
090924 11:31:19[ERROR]Fatal error:Can't open and lock privilege tables:Table'mysql.host'doesn't exist
```

```
090924 11:31:19 mysqld ended
```

```
mysql
warning
InnoDB
redo log
```

```
090924 11:39:44 InnoDB:ERROR:the age of the last checkpoint is 9433712,
```

```
InnoDB:which exceeds the log group capacity 9433498.
```

```
InnoDB:If you are using big BLOB or TEXT rows,you must set the
```

```
InnoDB:combined size of log files at least 10 times bigger than the
```

```
InnoDB:largest such row.
```

```
090924 11:40:00 InnoDB:ERROR:the age of the last checkpoint is 9433823,
```

```
InnoDB:which exceeds the log group capacity 9433498.
```

```
InnoDB:If you are using big BLOB or TEXT rows,you must set the
```

```
InnoDB:combined size of log files at least 10 times bigger than the
```

```
InnoDB:largest such row.
```

```
090924 11:40:16 InnoDB:ERROR:the age of the last checkpoint is 9433645,
```

```
InnoDB:which exceeds the log group capacity 9433498.
```

```
InnoDB:If you are using big BLOB or TEXT rows,you must set the
```

```
InnoDB:combined size of log files at least 10 times bigger than the
```

```
InnoDB:largest such row.
```

3.2.2 配置

3.2.1 配置慢查询日志slow log
DBA 需要配置 SQL 日志，MySQL 提供了两种 SQL 日志，一种是 DBA 常用的 long_query_time 配置，记录执行时间超过 10 秒的 SQL。

MySQL 配置 long_query_time 为 ON

```
mysql> SHOW VARIABLES LIKE 'long_query_time'\G;

*****1. row*****

Variable_name: long_query_time

Value: 10.000000

1 row in set (0.00 sec)

mysql> SHOW VARIABLES LIKE 'log_slow_queries'\G;

*****1. row*****

Variable_name: log_slow_queries

Value: ON

1 row in set (0.00 sec)
```

long_query_time 是 MySQL 配置 SQL 日志的变量，记录执行时间超过 long_query_time 的 SQL。MySQL 5.1 版本 long_query_time 默认值为 10 秒，DBA 可以根据需要调整，例如设置为 0.5 秒或 0.05 秒。

MySQL 配置 log_queries_not_using_indexes 为 ON，记录没有使用索引的 SQL。

```
mysql[SHOW VARIABLES LIKE 'log_queries_not_using_indexes'\G;
```

```
*****1.row*****
```

```
Variable_name:log_queries_not_using_indexes
```

```
Value:ON
```

```
1 row in set(0.00 sec)
```

MySQL 5.6.5

log_throttle_queries_not_using_indexes slow
log SQL 0
SQL slow log slow log DBA
DBA

DBA SQL MySQL
SQL
MySQL mysqldumpslog DBA
DBA

```
[root@nh122-190 data]#mysqldumpslow nh122-190-slow.log
```

```
Reading mysql slow query log from nh122-190-slow.log
```

```
Count:11 Time=10.00s(110s)Lock=0.00s(0s)Rows=0.0(0),dbother[dbother]@localhost
```

```
insert into test.DbStatus select now(),(N-com_select)/(N-uptime),(N-com_insert)/(N-uptime),(N-com_update)/(N-uptime),(N-com_delete)/(N-uptime),N-(N/N),N-(N/N),N.N/N,N-N/(N*N),GetCPULoadInfo(N)from test.CheckDbStatus order by check_id desc limit N
```

```
Count:653 Time=0.00s(0s)Lock=0.00s(0s)Rows=0.0(0),9Y0Ugs_SC[9Y0Ugs_SC]@[192.168.43.7]
```

```
select custom_name_one from'low_game_schema'.'role_details'where role_id='S'rse and summarize the MySQL slow query log.Options are
```

```
--verbose verbose
```

```
--debug debug
```

```
--help write this text to standard output
```

```
-v verbose
```

```
-d debug
```

```
-s ORDER what to sort by(al,at,ar,c,l,r,t),'at'is default
```

```
al:average lock time
```

```
ar:average rows sent
```

at:average query time

c:count

l:lock time

r:rows sent

t:query time

-r reverse the sort order(largest last instead of first)

-t NUM just show the top n queries

-a don't abstract all numbers to N and strings to 'S'

-n NUM abstract numbers with at least n digits within names

-g PATTERN grep:only consider stmts that include this string

-h HOSTNAME hostname of db server for*-slow.log filename(can be wildcard),
default is'',i.e.match all

-i NAME name of server instance(if using mysql.server startup script)

-l don't subtract lock time from total time

10 SQL

```
[root@nh119-141 data]#mysqldumpslow-s al-n 10 david.log

Reading mysql slow query log from david.log

Count:5 Time=0.00s(0s)Lock=0.20s(1s)Rows=4.4(22),Audition[Audition]@[192.168.30.108]

SELECT OtherSN,State FROM wait_friend_info WHERE UserSN=N

Count:1 Time=0.00s(0s)Lock=0.00s(0s)Rows=1.0(1),audition-kr[audition-kr]@[192.168.30.105]

SELECT COUNT(N)FROM famverifycode WHERE UserSN=N AND verifycode='S'

.....
```

MySQL 5.1

mysql slow_log

```
mysql[SHOW CREATE TABLE mysql.slow_log\G;

*****1.row*****
```

Table:slow_log

```
Create Table:CREATE TABLE'slow_log'()
```

```
'start_time'timestamp NOT NULL DEFAULT CURRENT_TIMESTAMP ON UPDATE CURRENT_TIMESTAMP,
```

```
'user_host'mediumtext NOT NULL,
```

```
'query_time'time NOT NULL,
```

```
'lock_time'time NOT NULL,
```

```
'rows_sent' int(11) NOT NULL,
```

```
'rows_examined' int(11) NOT NULL,
```

```
'db' varchar(512) NOT NULL,
```

```
'last_insert_id' int(11) NOT NULL,
```

```
'insert_id' int(11) NOT NULL,
```

```
'server_id' int(11) NOT NULL,
```

```
'sql_text' mediumtext NOT NULL
```

```
)ENGINE=CSV DEFAULT CHARSET=utf8 COMMENT='Slow log'
```

```
1 row in set(0.00 sec)
```

```
log_output=FILETABLE
mysqlslow_log
```

```
mysql> SHOW VARIABLES LIKE 'log_output'\G;
```

+ - - - - + - - - - +

| Variable_name | Value |
|---------------|-------|
|---------------|-------|

⊕ - - - - - ⊕ - - - - - ⊕

```
|log_output|FILE|
```

+

```
1 row in set(0.00 sec)
```

```
mysql> SET GLOBAL log_output='TABLE';
```

Query OK, 0 rows affected (0.00 sec)

```
mysql> SHOW VARIABLES LIKE 'log_output'\G;
```

+ - - - - + - - - - +

|Variable_name|Value|

+-----+-----+

|log_output|TABLE|

+-----+-----+

1 row in set(0.00 sec)

mysql>select sleep(10)\G;

+-----+

|sleep(10)|

+-----+

|0|

+-----+

1 row in set(10.01 sec)

mysql>SELECT*FROM mysql.slow_log\G;

*****1. row*****

start_time:2009-09-25 13:44:29

user_host:david[david]@localhost[]

query_time:00:00:09

lock_time:00:00:00

rows_sent:1

rows_examined:0

db:mysql

last_insert_id:0

insert_id:0

server_id:0

sql_text:select sleep(10)

1 row in set(0.00 sec)

log_output
sleep10SQLslow_log

slow_log CSV
slow_log MyISAM start_time
MySQL

```
mysql> ALTER TABLE mysql.slow_log ENGINE=MyISM;
```

```
ERROR 1580 (HY000): You cannot 'ALTER' a log table if logging is enabled
```

```
mysql> SET GLOBAL slow_query_log=off;
```

```
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> ALTER TABLE mysql.slow_log ENGINE=MyISAM;
```

```
Query OK, 1 row affected (0.00 sec)
```

```
Records: 1 Duplicates: 0 Warnings: 0
```

slow_log MyISAM
MySQL

MySQL slow log SQL
SQL
0.5

InnoDB SQL MySQL slow log
logical reads physical reads
IO

```
#Time:111227 23:49:16
```

```
#User@Host:root[root]@localhost[127.0.0.1]
```

```
#Query_time:6.081214 Lock_time:0.046800 Rows_sent:42 Rows_examined:727558 Logical_reads:91584 Physical_reads:19
```

```
use tpcc;
```

```
SET timestamp=1325000956;
```

```
SELECT orderid,customerid,employeeid,orderdate
```

```
FROM orders
```

```
WHERE orderdate IN
```

```
(SELECT MAX(orderdate)
```

```
FROM orders
```

```
GROUP BY (DATE_FORMAT(orderdate, '%Y%M'))
```

```
);
```

91 58419SQL

long_query_ioIOSQLslow log
100100SQLslow log
MySQLslow_query_typeslow log

0SQLslow log

1SQLslow log

2IOSQLslow log

3IOSQLslow log

3.2.3 配置

MySQL 的日志文件在 `/var/log/mysql/` 目录下，默认情况下，MySQL 会记录所有的连接和查询操作。

```
[root@nineyou0-43 data]#tail nineyou0-43.log

090925 11:00:24 44 Connect zlm@192.168.0.100 on

44 Query SET AUTOCOMMIT=0

44 Query set autocommit=0

44 Quit

090925 11:02:37 45 Connect Access denied for user'root'@'localhost'(using password:N0)

090925 11:03:51 46 Connect Access denied for user'root'@'localhost'(using password:N0)

090925 11:04:38 23 Query rollback
```

Access denied 表示用户没有权限执行操作。MySQL 5.1 版本开始，MySQL 提供了 `general_log` 和 `slow_log` 两个日志文件，用于记录所有的连接和查询操作。

3.2.4 配置

MySQL 的 binary log 记录 MySQL 的 SQL 语句。例如，SELECT 和 SHOW 语句。MySQL 的 binary log 记录 MySQL 的 SQL 语句。例如，SELECT 和 SHOW 语句。

```
mysql>UPDATE t SET a=1 WHERE a=2;
```

```
Query OK, 0 rows affected (0.00 sec)
```

```
Rows matched: 0 Changed: 0 Warnings: 0
```

```
mysql>SHOW MASTER STATUS\G;
```

```
*****1.row*****
```

```
File:mysql_d.000008
```

```
Position:383
```

```
Binlog_Do_DB:
```

```
Binlog_Ignore_DB:
```

```
Executed_Gtid_Set:
```

```
1 row in set (0.00 sec)
```

```
mysql>SHOW BINLOG EVENTS IN 'mysql_d.000008'\G;
```

```
*****1.row*****
```

```
Log_name:mysql_d.000008
```

```
Pos:4
```

```
Event_type:Format_desc
```

```
Server_id:1
```

```
End_log_pos:120
```

```
Info:Server ver:5.6.6-m9-log,Binlog ver:4
```

```
*****2.row*****
```

```
Log_name:mysql_d.000008
```

```
Pos:120
```

```
Event_type:Query
```

```
Server_id:1
```

End_log_pos:199

Info:BEGIN

*****3.row*****

Log_name:mysql.000008

Pos:199

Event_type:Query

Server_id:1

End_log_pos:303

Info:use'test';UPDATE t SET a=1 WHERE a=2

*****4.row*****

Log_name:mysql.000008

Pos:303

Event_type:Query

Server_id:1

End_log_pos:383

Info:COMMIT

4 rows in set(0.00 sec)

MySQL UPDATE Changed
0 SHOW BINLOG EVENT

SELECT SHOW

recovery
point-in-time

replication MySQL
slave standby MySQL master primary

audit

mysql> show variables like 'log-bin[=name]';
+-----+-----+
|Variable_name|Value|
+-----+-----+
|datadir|usr/local/mysql/data|
+-----+-----+

mysql> show variables like 'datadir';

+-----+-----+

|Variable_name|Value|

+-----+-----+

|datadir|usr/local/mysql/data|

+-----+-----+

1 row in set (0.00 sec)

mysql> system ls -lh /usr/local/mysql/data/;

total 2.1G

-rw-rw---- 1 mysql mysql 6.5M Sep 25 15:13 bin_log.000001

-rw-rw---- 1 mysql mysql 17 Sep 25 00:32 bin_log.index

-rw-rw---- 1 mysql mysql 300M Sep 25 15:13 ibdata1

-rw-rw---- 1 mysql mysql 256M Sep 25 15:13 ib_logfile0

-rw-rw---- 1 mysql mysql 256M Sep 25 15:13 ib_logfile1

drwxr-xr-x 2 mysql mysql 4.0K May 7 10:08 mysql

drwx----- 2 mysql mysql 4.0K May 7 10:09 test

mysql> show binlog status;
+-----+-----+
|bin_log.000001|
|bin_log.index|
+-----+-----+

mysql> show variables like 'log_bin_log_size';
+-----+-----+
|Variable_name|Value|
+-----+-----+
|log_bin_log_size|100000000|
+-----+-----+

mysql> show variables like 'max_binlog_size';

mysql> show variables like 'max_binlog_size';

☐binlog_cache_size

☐sync_binlog

☐binlog-do-db

☐binlog-ignore-db

☐log-slave-update

☐binlog_format

max_binlog_size 1073741824
+1 MySQL 5.0 1 073 741 824
1 G max_binlog_size 1.1G

InnoDB uncommitted
committed
binlog_cache_size 32K
binlog_cache_size session
MySQL binlog_cache_size
binlog_cache_size MySQL
SHOW GLOBAL STATUS
binlog_cache_use binlog_cache_disk_use
binlog_cache_size Binlog_cache_use
binlog_cache_disk_use

```
mysql>show variables like'binlog_cache_size';
```

```
+-----+-----+
```

```
|Variable_name|Value|
```

```
+-----+-----+
```

```
|binlog_cache_size|32768|
```

```
+-----+-----+
```

```
1 row in set(0.00 sec)
```



```
mysql>show global status like'binlog_cache%';
```

```
+-----+-----+
|Variable_name|Value|
+-----+-----+
|binlog_cache_disk_use|0|
|binlog_cache_use|33553|
+-----+-----+

2 rows in set(0.00 sec)
```

可以看到33 553个二进制日志缓存0个32KB二进制日志缓存MySQL
二进制日志缓存binlog_cache_size

二进制日志缓存二进制日志缓存二进制日志缓存二进制日志缓存二进制日志缓存
二进制日志缓存二进制日志缓存二进制日志缓存二进制日志缓存二进制日志缓存
sync_binlog=[N]二进制日志缓存二进制日志缓存N1sync_binlog=1
二进制日志缓存二进制日志缓存二进制日志缓存二进制日志缓存二进制日志缓存
sync_binlog0InnoDB二进制日志缓存二进制日志缓存二进制日志缓存
二进制日志缓存ON二进制日志缓存ON二进制日志缓存IO二进制日志缓存

二进制日志缓存sync_binlog1二进制日志缓存二进制日志缓存InnoDB二进制日志缓存
二进制日志缓存COMMIT二进制日志缓存sync_binlog1二进制日志缓存二进制日志缓存
二进制日志缓存二进制日志缓存二进制日志缓存二进制日志缓存二进制日志缓存MySQL二进制日志缓存
二进制日志缓存COMMIT二进制日志缓存二进制日志缓存二进制日志缓存二进制日志缓存二进制日志缓存
二进制日志缓存二进制日志缓存innodb_support_xa1二进制日志缓存
innodb_support_xaXA二进制日志缓存二进制日志缓存二进制日志缓存InnoDB二进制日志缓存
二进制日志缓存

二进制日志缓存binlog-do-db二进制日志缓存ignore-db二进制日志缓存二进制日志缓存二进制日志缓存
二进制日志缓存二进制日志缓存二进制日志缓存

二进制日志缓存slave二进制日志缓存master二进制日志缓存二进制日志缓存二进制日志缓存
二进制日志缓存二进制日志缓存log-slave-update二进制日志缓存master=slave=
slave二进制日志缓存二进制日志缓存二进制日志缓存

binlog_format 是 MySQL 5.1 中引入的一个新选项，用于控制二进制日志的格式。它支持以下几种格式：
SQL statement 格式：记录完整的 SQL 语句。
Replication 格式：记录 Oracle Standby 所需的复制信息。
rand uuid 格式：记录随机 UUID。
not sync 格式：记录 InnoDB 的同步信息。
REPEATABLE READ 格式：记录可重复读的快照。
COMMITTED 格式：记录 Oracle 和 Microsoft SQL Server 的提交信息。

MySQL 5.1 中 binlog_format 的默认值是 STATEMENT 和 ROW 的混合，即 MIXED。

1. STATEMENT 格式：MySQL 记录完整的 SQL 语句。

2. ROW 格式：MySQL 记录行级别的更改。ROW 格式是 Oracle Standby 所需的。Statement 格式是 MySQL 5.1 中的默认格式。binlog_format 为 ROW 时，InnoDB 的默认读模式为 READ COMMITTED。

3. MIXED 格式：MySQL 记录 STATEMENT 和 ROW 的混合。

1. 在 NDB 中，DML 语句记录为 ROW 格式。

2. 在 UUID()、USER()、CURRENT_USER()、FOUND_ROWS()、ROW_COUNT() 等函数中。

3. 在 INSERT DELAY 语句中。

4. 在 UDF 函数中。

5. 在 temporary table 中。

binlog_format 的默认值是 3-1。

表 3-1 存储引擎对二进制日志格式的支持情况

| 存储引擎 | Row 格式 | Statement 格式 |
|-----------|--------|--------------|
| InnoDB | Yes | Yes |
| MyISAM | Yes | Yes |
| HEAP | Yes | Yes |
| MERGE | Yes | Yes |
| NDB | Yes | No |
| Archive | Yes | Yes |
| CSV | Yes | Yes |
| Federate | Yes | Yes |
| Blockhole | No | Yes |

binlog_format=ROW;
binlog_format=ROW;

```
mysql>SET@session.binlog_format='ROW';  
  
Query OK,0 rows affected(0.00 sec)  
  
mysql>SELECT@session.binlog_format;
```

```
+-----+
|@@session.binlog_format|
+-----+
|ROW|
+-----+

1 row in set(0.00 sec)
```

binlog_format

```
mysql>SET GLOBAL binlog_format='ROW';

Query OK,0 rows affected(0.00 sec)

mysql>SELECT@@global.binlog_format;
+-----+
|@@global.binlog_format|
+-----+
|ROW|
+-----+

1 row in set(0.00 sec)
```

binlog_format ROW

ROW

100W UPDATE

```
mysql>SELECT@@session.binlog_format\G;

*****1.row*****

@@session.binlog_format:STATEMENT

1 row in set(0.00 sec)

mysql>SHOW MASTER STATUS\G;

*****1.row*****
```

File:test.000003

Position:106

Binlog_Do_DB:

Binlog_Ignore_DB:

1 row in set(0.00 sec)

mysql>UPDATE t1 SET username=UPPER(username);

Query OK,89279 rows affected(1.83 sec)

Rows matched:100000 Changed:89279 Warnings:0

mysql>SHOW MASTER STATUS\G;

*****1.row*****

File:test.000003

Position:306

Binlog_Do_DB:

Binlog_Ignore_DB:

1 row in set(0.00 sec)

binlog_format=STATEMENTUPDATE
200306-106ROWt2

mysql>SET SESSION binlog_format='ROW';

Query OK,0 rows affected(0.00 sec)

mysql>SHOW MASTER STATUS\G;

*****1.row*****

File:test.000003

Position:306

Binlog_Do_DB:

Binlog_Ignore_DB:

1 row in set(0.00 sec)

mysql>UPDATE t2 SET username=UPPER(username);

Query OK,89279 rows affected(2.42 sec)

Rows matched:100000 Changed:89279 Warnings:0

mysql>SHOW MASTER STATUS\G;

*****1.row*****

File:test.000003

Position:13782400

Binlog_Do_DB:

Binlog_Ignore_DB:

1 row in set(0.00 sec)

ROW13 782 094
13MBt217MB
sync_binlog=1MySQLSQL

binlog_formatROW

cat
headtailMySQL
mysqlbinlogSTATEMENTmysqlbinlog
SQL

[root@nineyou0-43 data]#mysqlbinlog--start-position=203 test.000004

/*!40019 SET@session.max_insert_delayed_threads=0*/;

...

#090927 15:43:11 server id 1 end_log_pos 376 Query thread_id=188 exec_time=1 error_code=0

SET TIMESTAMP=1254037391/*!*/;

update t2 set username=upper(username)where id=1

/*!*/;

#at 376

#090927 15:43:11 server id 1 end_log_pos 403 Xid=1009

COMMIT/*!*/;

```
DELIMITER;

#End of log file

ROLLBACK/*added by mysqlbinlog*/;

/*!50003 SET COMPLETION_TYPE=@OLD_COMPLETION_TYPE*/;
```

```
SQLUPDATE t2 SET username=UPPERusername
WHERE id=1SQL
mysqlbinlogOracle LogMinerROW
mysqlbinlog“”unreadable
```

```
[root@nineyou0-43 data]#mysqlbinlog--start-position=1065 test.000004

/*!40019 SET@session.max_insert_delayed_threads=0*/;

.....

#at 1135

#at 1198

#090927 15:53:52 server id 1 end_log_pos 1198 Table_map:'member'.'t2'mapped to number 58

#090927 15:53:52 server id 1 end_log_pos 1378 Update_rows:table id 58 flags:STMT_END_F

BINLOG'

EBq/ShMBAAAPwAAAK4EAAAAADoAAAAAABm1lbWJlcgACdDIACgMPDw/+CgsPAQwKJAAoAEAA

/gJAAAAA

EBq/ShgBAAAAtAAAAAGIFAAQADoAAAAAAEACv///8A/AEAAALYwx1eDk50Dh5b3UE0X1vdSA3

Y2JiMzI1MmJhNmI3ZTljNDIyZmFjNTMzNGQyMjA1NAFNLacPAAAAABjEnpxPBIAAD8AQAAAAtB

TEVY0Tk40F1PVQ05eW91IDdjYmIzMjUyYmE2Yjd1OWM0MjJmYW1MzM0ZDIyMDU0AU0tpw8AAAAA

AGMSenE8EgAA

'/*!*/;

#at 1378

#090927 15:53:52 server id 1 end_log_pos 1405 Xid=1110

COMMIT/*!*/;

DELIMITER;

#End of log file

ROLLBACK/*added by mysqlbinlog*/;
```

```
/*!50003 SET COMPLETION_TYPE=@OLD_COMPLETION_TYPE*/;
```

SQL-v-vv
-vv-v-vv

```
[root@nineyou0-43 data]#mysqlbinlog-vv--start-position=1065 test.000004
```

.....

```
BINLOG'
```

```
EBq/ShMBAAAApAAAK4EAAAAADoAAAAAABm1lbWJlcgACdDIACgMPDw/+CgsPAQwKJAAoEAA
```

```
/gJAAAAA
```

```
EBq/ShgBAAAAaAAAGIFAAQADoAAAAAAEACv///8A/AEAAAALYwxleDk50Dh5b3UE0XlvdSA3
```

```
Y2JiMzI1MmJhNmI3ZTljNDIyZmFjNTMzNGQyMjA1NAFNLaCpAAAAABjEnpxPBIAAAD8AQAAAAtB
```

```
TEVY0Tk40FlPVQ05eW9lIDdjYmIzMjUyYmE2YjdLOWM0MjJmYW1MzM0ZDIyMDU0AU0tpw8AAAAA
```

```
AGMSenE8EgAA
```

```
'/*!*/;
```

```
###UPDATE member.t2
```

```
###WHERE
```

```
###@1=1/*INT meta=0 nullable=0 is_null=0*/
```

```
###@2='david'/*VARSTRING(36)meta=36 nullable=0 is_null=0*/
```

```
###@3='family'/*VARSTRING(40)meta=40 nullable=0 is_null=0*/
```

```
###@4='7cbb3252ba6b7e9c422fac5334d22054'/*VARSTRING(64)meta=64 nullable=0 is_null=0*/
```

```
###@5='M'/*STRING(2)meta=65026 nullable=0 is_null=0*/
```

```
###@6='2009:09:13'/*DATE meta=0 nullable=0 is_null=0*/
```

```
###@7='00:00:00'/*TIME meta=0 nullable=0 is_null=0*/
```

```
###@8=''/*VARSTRING(64)meta=64 nullable=0 is_null=0*/
```

```
###@9=0/*TINYINT meta=0 nullable=0 is_null=0*/
```

```
###@10=2009-08-11 16:32:35/*DATETIME meta=0 nullable=0 is_null=0*/
```

```
###SET
```

```
###@1=1/*INT meta=0 nullable=0 is_null=0*/
```

```
###@2='DAVID'/*VARSTRING(36)meta=36 nullable=0 is_null=0*/
```



```
###@3=family/*VARSTRING(40)meta=40 nullable=0 is_null=0*/

###@4='7cbb3252ba6b7e9c422fac5334d22054'/*VARSTRING(64)meta=64 nullable=0 is_null=0*/

###@5='M'/*STRING(2)meta=65026 nullable=0 is_null=0*/

###@6='2009:09:13'/*DATE meta=0 nullable=0 is_null=0*/

###@7='00:00:00'/*TIME meta=0 nullable=0 is_null=0*/

###@8=''/*VARSTRING(64)meta=64 nullable=0 is_null=0*/

###@9=0/*TINYINT meta=0 nullable=0 is_null=0*/

###@10=2009-08-11 16:32:35/*DATETIME meta=0 nullable=0 is_null=0*/

#at 1378

#090927 15:53:52 server id 1 end_log_pos 1405 Xid=1110

COMMIT/*!*/;

DELIMITER;

#End of log file

ROLLBACK/*added by mysqlbinlog*/;

/*!50003 SET COMPLETION_TYPE=@OLD_COMPLETION_TYPE*/;
```

mysqlbinlog 更新 t2 设置
username=upper(username)where id=1
10W ROW 13MB

3.3 設定

MySQLはUNIX系OSで動作するため、MySQLもUNIX系OSで動作するsocketを使用する。MySQLのsocketは/tmp/mysql.sockで指定する。

```
mysql> SHOW VARIABLES LIKE 'socket'\G;
```

```
*****1.row*****
```

```
Variable_name: socket
```

```
Value: /tmp/mysql.sock
```

```
1 row in set (0.00 sec)
```

3.4 pid

MySQLのID——pid
pid_file

```
mysql>show variables like'pid_file'\G;
```

```
*****1.row*****
```

```
Variable_name:pid_file
```

```
Value:/usr/local/mysql/data/xen-server.pid
```

```
1 row in set(0.00 sec)
```

3.5 □□□□□□□

```
MySQLMySQLMySQLfrm
```

```
frm v_a v_a.frm
cat
```

```
[root@xen-server test]#cat v_a.frm
```

TYPE=VIEW

```
query=select 'test'. 'a'. 'b' AS 'b' from 'test'. 'a'
```

md5=4eda70387716a4d6c96f3042dd68b742

updatable=1

```
algorithm=0
```

```
definer_user=root
```

```
definer_host=localhost
```

suid=2

```
with_check_option=0
```

timestamp=2010-08-04 07:23:36

```
create-version=1
```

```
source=select*from a
```

```
client_cs_name=utf8
```

```
connection_cl_name=utf8_general_ci
```

```
view_body_utf8=select'test'.'a'.'b'AS'b'from'test'.'a'
```

3.6 InnoDB

MySQL InnoDB

3.6.1

InnoDB tablespace
10MB ibdata1 tablespace file
innodb_data_file_path

```
innodb_data_file_path=datafile_spec1[,datafile_spec2]...
```

```
[mysqld]
```

```
innodb_data_file_path=/db/ibdata1:2000M;/dr2/db/ibdata2:2000M:autoextend
```

/db/ibdata1/dr2/db/ibdata2
ibdata1 2000MB ibdata2 2000MB
2000MB autoextend

innodb_data_file_path InnoDB
innodb_file_per_table InnoDB
.ibd
MySQL innodb_file_per_table

```
mysql> SHOW VARIABLES LIKE 'innodb_file_per_table';
```

```
*****1.row*****
```

```
Variable_name:innodb_file_per_table
```

Value:0N

1 row in set(0.00 sec)

mysql\system ls-lh/usr/local/mysql/data/member/*

-rw-r-----1 mysql mysql 8.7K 2009-02-24/usr/local/mysql/data/member/Profile.frm

-rw-r-----1 mysql mysql 1.7G 9 11:13/usr/local/mysql/data/member/Profile.ibd

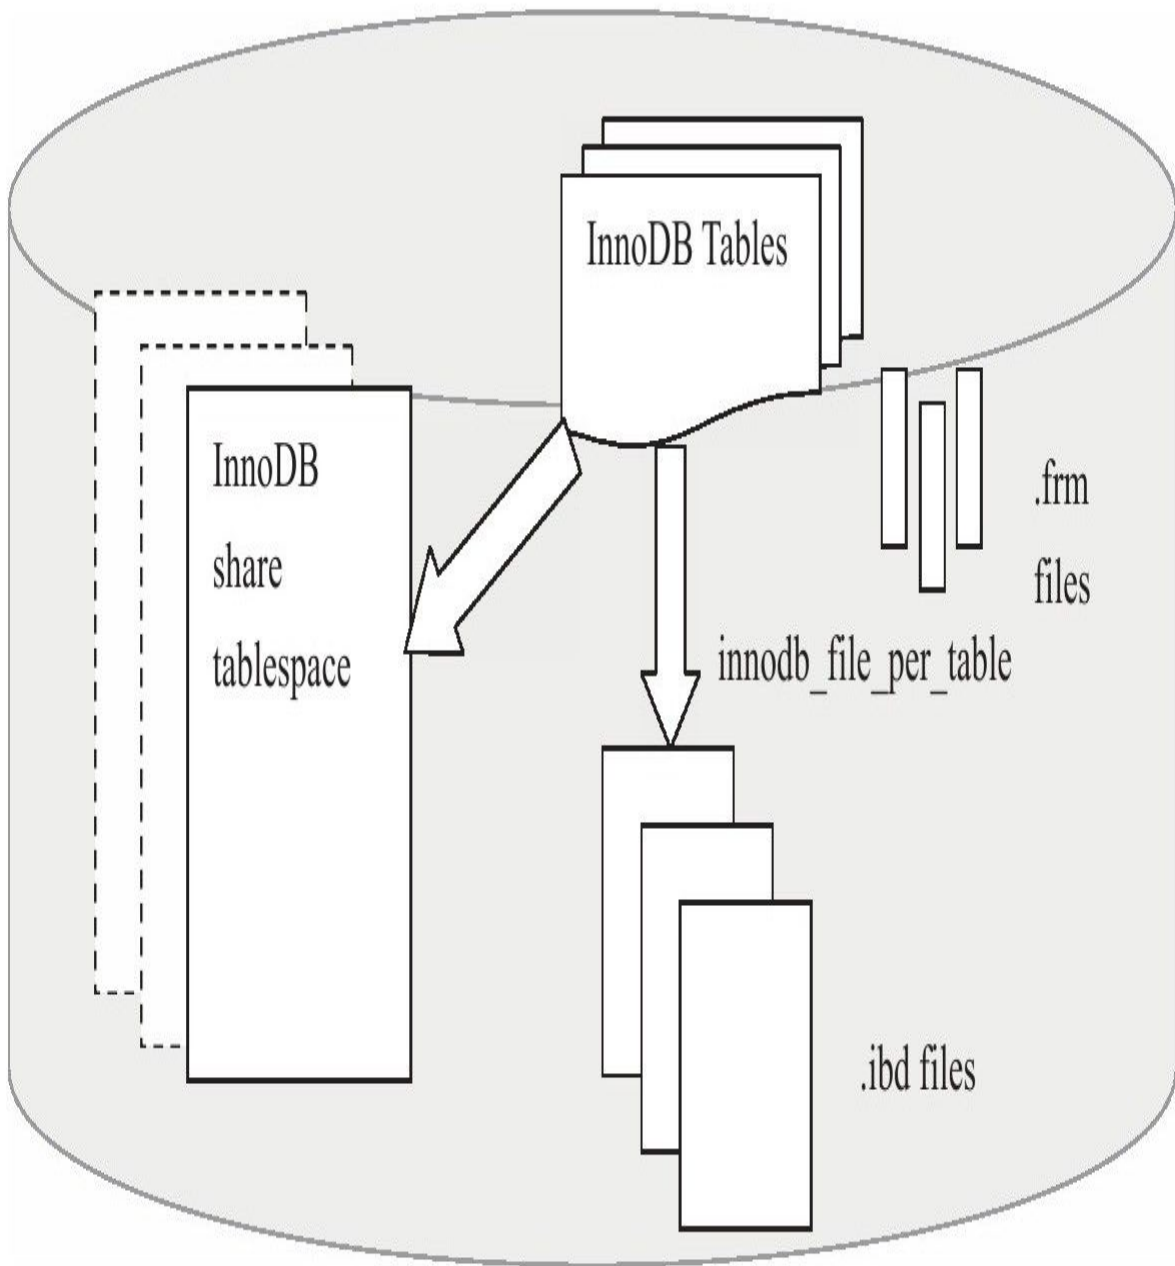
-rw-rw----1 mysql mysql 8.7K 9 13:38/usr/local/mysql/data/member/t1.frm

-rw-rw----1 mysql mysql 17M 9 13:40/usr/local/mysql/data/member/t1.ibd

-rw-rw----1 mysql mysql 8.7K 9 15:42/usr/local/mysql/data/member/t2.frm

-rw-rw----1 mysql mysql 17M 9 15:54/usr/local/mysql/data/member/t2.ibd

Profile t1 t2 InnoDB
innodb_file_per_table=ON .ibd
BITMAP
3-1 InnoDB



□ 3-1 InnoDB□□□□□□

3.6.2 五五五五五

```

#####InnoDB#####ib_logfile0ib_logfile1
#####MySQL#####InnoDB#####
#####redo log file#####InnoDB#####
#####InnoDB#####

```

```

media failure
InnoDB

```

```

InnoDB1group2
ib_logfile0ib_logfile1
mirrored log groups
InnoDB1
221
3-23

```

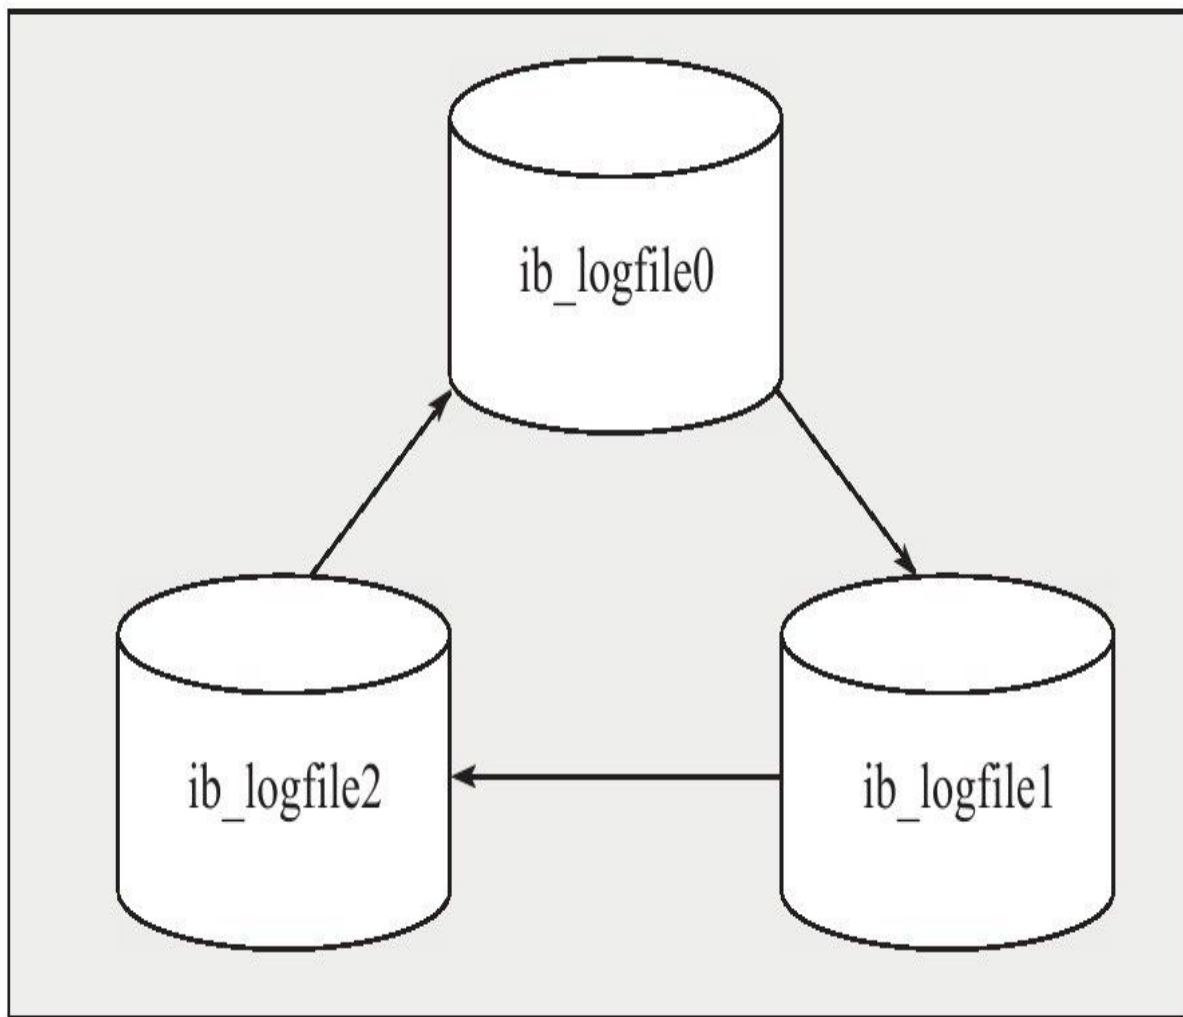



图 3-2 InnoDB 日志文件

配置参数如下：

- innodb_log_file_size
- innodb_log_files_in_group
- innodb_mirrored_log_groups
- innodb_log_group_home_dir

innodb_log_file_size 设置 InnoDB 1.2.x 版本
4GB 1.2.x 版本 512GB

innodb_log_files_in_group 设置 2
innodb_mirrored_log_groups 设置 1
innodb_log_group_home_dir 设置 ./
MySQL

```
mysql>SHOW VARIABLES LIKE'innodb%log%'\G;.....
```

```
*****4.row*****
```

```
Variable_name:innodb_log_file_size
```

```
Value:5242880
```

```
*****5.row*****
```

```
Variable_name:innodb_log_files_in_group
```

```
Value:2
```

```
*****6.row*****
```

```
Variable_name:innodb_log_group_home_dir
```

```
Value:./
```

```
*****7.row*****
```

```
Variable_name:innodb_mirrored_log_groups
```

```
Value:1
```

```
7 rows in set(0.00 sec)
```

InnoDB 版本 5.6.17 支持 async checkpoint
async checkpoint
InnoDB

```
090924 11:39:44 InnoDB:ERROR:the age of the last checkpoint is 9433712,
```

```
InnoDB:which exceeds the log group capacity 9433498.
```

```
InnoDB:If you are using big BLOB or TEXT rows,you must set the
```


表 3-2 重做日志条目结构

| redo_log_type | space | page_no | redo_log_body |
|---------------|-------|---------|---------------|
|---------------|-------|---------|---------------|

图 3-2 重做日志条目结构

redo_log_type 1 字节

space 4 字节 ID

page_no 4 字节

redo_log_body 可变长度

图 2 重做日志缓冲区结构

图 3-3 重做日志缓冲区结构

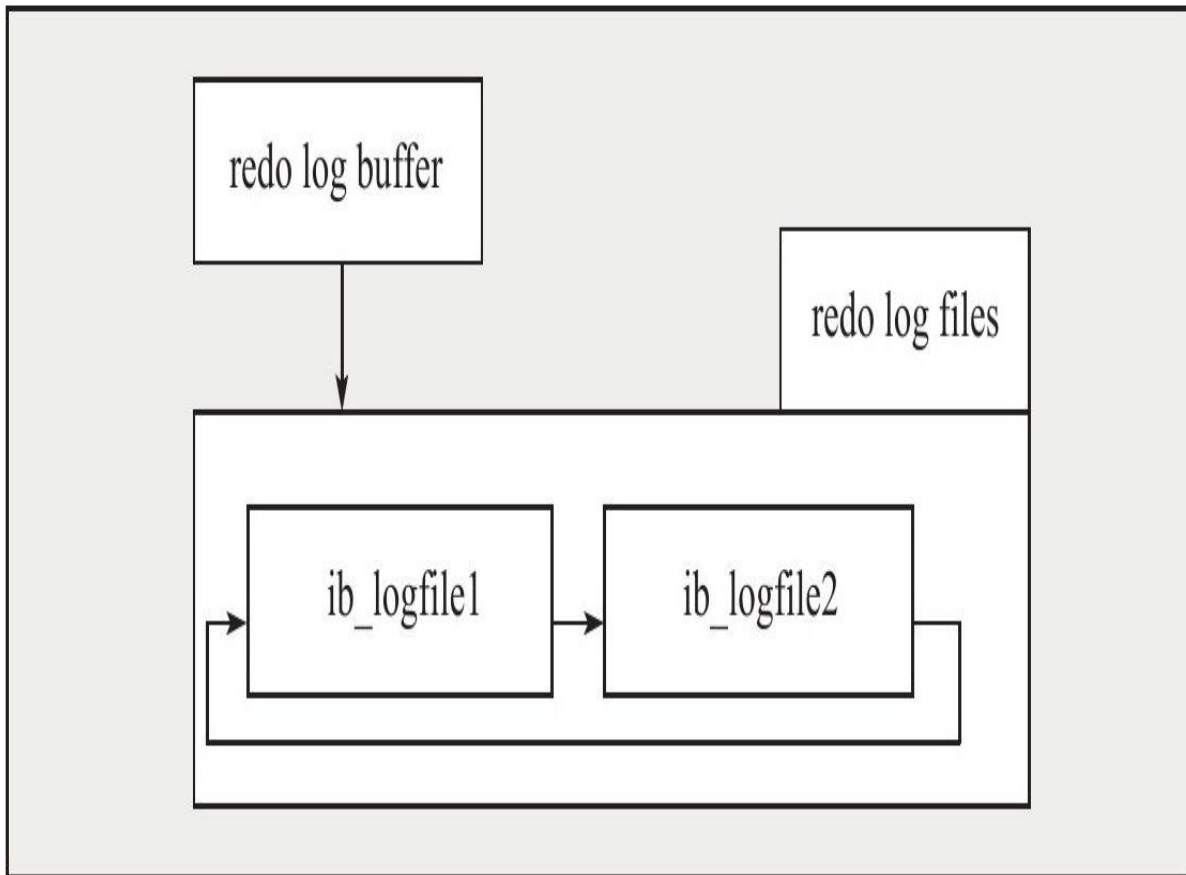


图 3-3 InnoDB 重做日志

InnoDB 的 redo log 缓冲区大小为 512 页，每页大小为 16KB。当 redo log 缓冲区写满时，InnoDB 会执行 doublewrite 操作。

当 redo log 缓冲区写满时，InnoDB 会执行 doublewrite 操作。此时，主线程（master thread）会调用 innodb_flush_log_at_trx_commit 函数，将 redo log 缓冲区中的数据 flush 到 redo log 文件中。该函数的参数为 1，表示在每次提交（commit）时 flush。

innodb_flush_log_at_trx_commit 函数的参数为 0、1 或 2。当参数为 0 时，表示在每次 flush 时 flush；当参数为 1 时，表示在每次提交（commit）时 flush；当参数为 2 时，表示在每次 flush 且提交（commit）时 flush。

MySQL ACID innodb_flush_log_at_trx_commit
1
0 2
2 MySQL

3.7 备份

MySQL 数据库的备份和恢复是数据库管理员（DBA）的重要任务之一。MySQL 数据库的备份和恢复可以通过多种方式进行，包括物理备份、逻辑备份和增量备份等。物理备份是指将数据库的物理文件（如数据文件、日志文件等）进行备份，而逻辑备份则是将数据库中的数据以 SQL 语句的形式进行备份。增量备份是指在已经进行了全量备份的基础上，只备份自上次备份以来发生变化的数据。

MySQL 数据库的备份和恢复还可以通过 point in time replication 来实现。point in time replication 是指在指定的时间点将数据库的状态复制到另一个数据库上。MySQL 5.1 版本引入了 STATEMENT 和 ROW 两种复制模式，其中 ROW 模式可以支持增量备份。DBA 可以根据需要选择合适的备份和恢复策略。

InnoDB 是 MySQL 数据库的默认存储引擎，它支持事务、行级锁和崩溃恢复等功能。InnoDB 数据库的备份和恢复可以通过多种方式进行，包括物理备份、逻辑备份和增量备份等。InnoDB 数据库的物理文件（如数据文件、日志文件等）可以通过物理备份的方式进行备份。InnoDB 数据库的逻辑备份可以通过 mysqldump 工具进行。

4

InnoDB
InnoDB

4.1 인덱스

InnoDB는 index organized table이다. InnoDB는 Primary Key를 인덱스화한다. InnoDB는

Unique NOT NULL를 인덱스화한다

InnoDB는 6를 인덱스화한다

InnoDB는 인덱스화한다

```
mysql>CREATE TABLE z(
-a INT NOT NULL,
-b INT NULL,
-c INT NOT NULL,
-d INT NOT NULL,
-UNIQUE KEY(b),
-UNIQUE KEY(d),UNIQUE KEY(c));
Query OK,0 rows affected(0.02 sec)

mysql>INSERT INTO z SELECT 1,2,3,4;
Query OK,1 row affected(0.00 sec)

Records:1 Duplicates:0 Warnings:0

mysql>INSERT INTO z SELECT 5,6,7,8;
Query OK,1 row affected(0.00 sec)

Records:1 Duplicates:0 Warnings:0

mysql>INSERT INTO z SELECT 9,10,11,12;
Query OK,1 row affected(0.00 sec)

Records:1 Duplicates:0 Warnings:0
```

```

0000000000z00a0b0c0d0000b0c0d0000000000000000b0000
NULL000000000000000000000000000000000000SQL00000000
00

```

```
mysql> SELECT a,b,c,d,_rowid FROM z;
```

+ - - - + - - - - - + - - - - + - - - - + - - - - - - - - +

```
|a|b|c|d|_rowid|
```

+-+-+-----+-----+-----+-----+-----+

|1|2|3|4|4|

|5|6|7|8|8|

| 9 | 10 | 11 | 12 | 12 |

+-+-+-----+-----+-----+-----+-----+

```
3 rows in set (0.00 sec)
```

```
_rowid zc d
d InnoDB

```

rowid

```
mysql>CREATE TABLE a(
```

-[]a INT,

-[]b INT,

- PRIMARY KEY(a,b)

- □) ENGINE=InnoDB;

Query OK, 0 rows affected (0.03 sec)

```
mysql>INSERT INTO a SELECT 1,1;
```

Query OK, 1 row affected (0.01 sec)

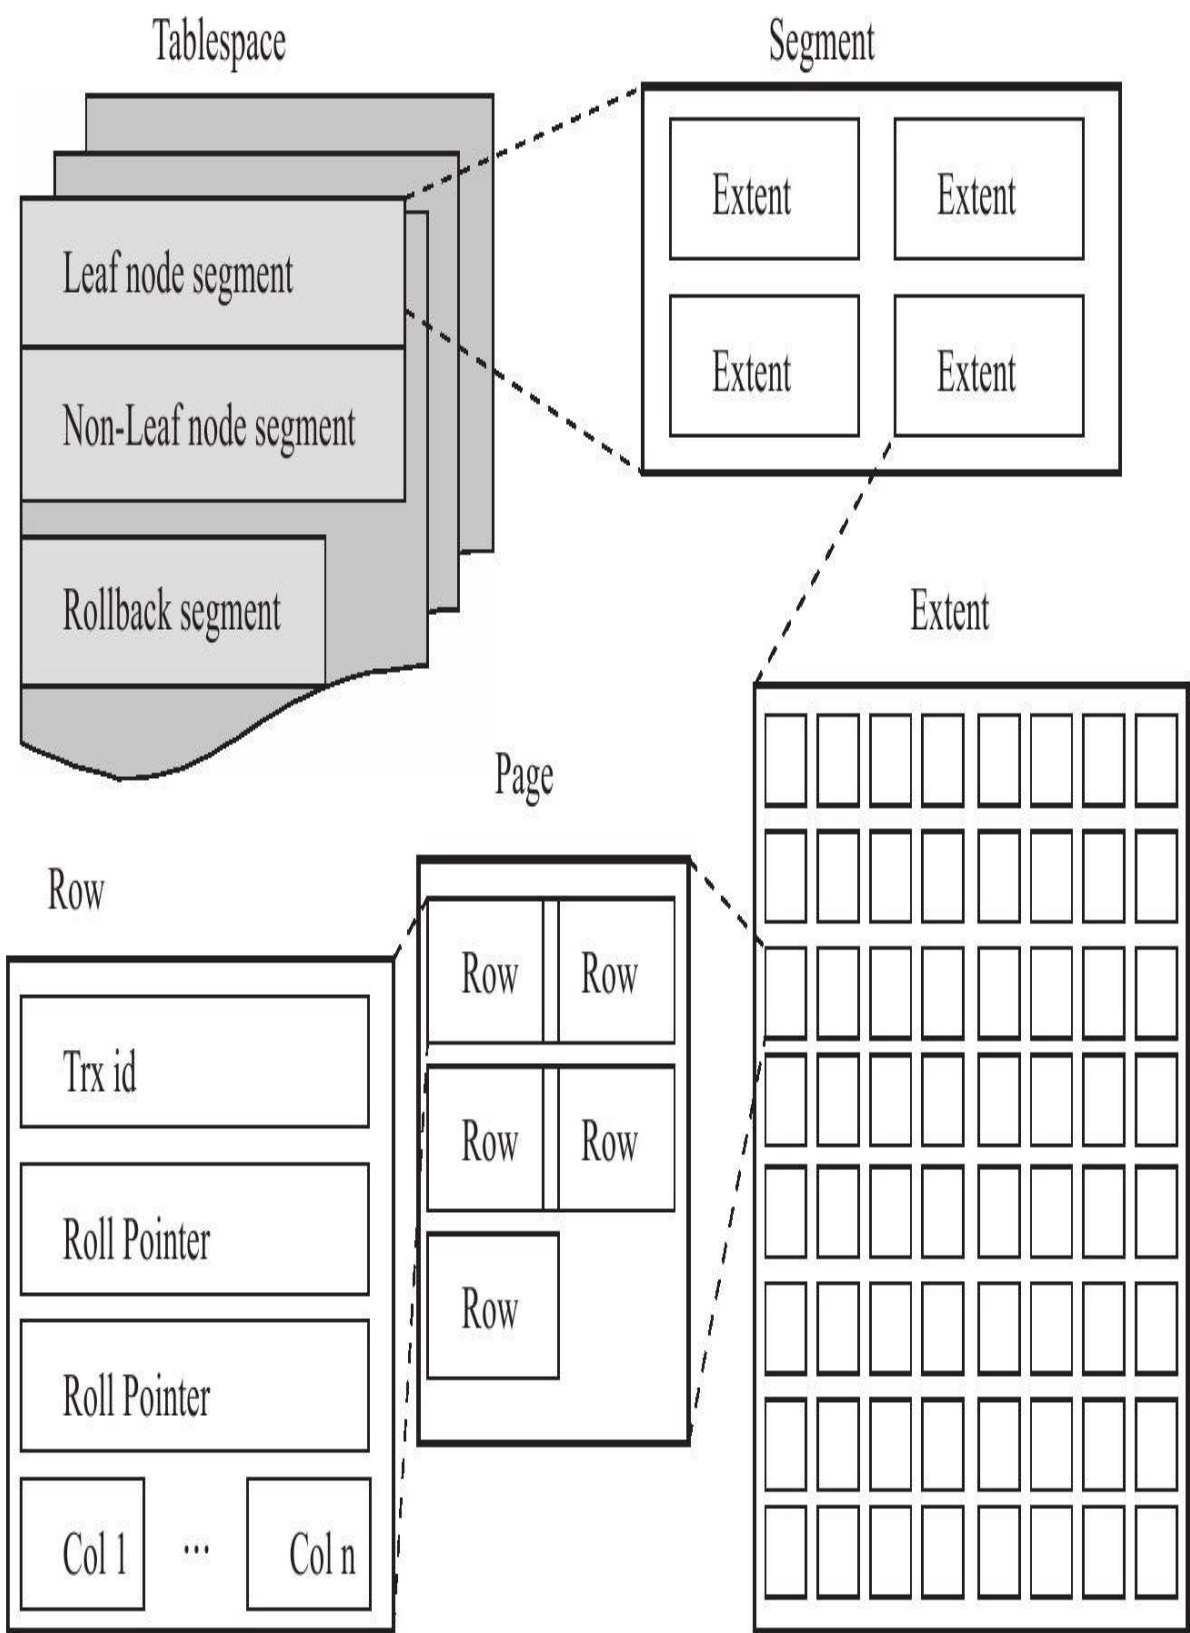
Records:1 Duplicates:0 Warnings:0

```
mysql> SELECT a, _rowid FROM a;
```

ERROR 1054(42S22):Unknown column '_rowid' in 'field list'

4.2 InnoDB

InnoDB
tablespace segment extent page
block InnoDB 4-1



4-1 InnoDB

4.2.1

InnoDB 3
InnoDB ibdata1
innodb_file_per_table

innodb_file_per_table
Bitmap undo
Double write buffer
innodb_file_per_table
innodb_file_per_table ON

```
mysql>SHOW VARIABLES LIKE'innodb_file_per_table'\G;
```

```
*****1.row*****
```

```
Variable_name:innodb_file_per_table
```

```
Value:ON
```

```
1 row in set(0.00 sec)
```

```
mysql>system ls-lh/usr/local/mysql/data/ibdata*
```

```
-rw-rw----1 mysql mysql 58M Mar 11 13:58/usr/local/mysql/data/ibdata1
```

ibdata1 58MB undo 1
mytest InnoDB

```
mysql>SET autocommit=0;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>UPDATE mytest SET salary=0;
```

```
Query OK,2844047 rows affected(19.47 sec)
```

```
Rows matched:2844047 Changed:2844047 Warnings:0
```

```
mysql>system ls-lh/usr/local/mysql/data/ibdata*
```

```
-rw-rw----1 mysql mysql 114M Mar 11 14:00/usr/local/mysql/data/ibdata1
```

0
commit rollback undo update mytest set
salary=0 ibdata1 114MB
undo

k rollback ibdata1
58MB

```
mysql> ROLLBACK;

Query OK, 0 rows affected (0.00 sec)

mysql> system ls -lh /usr/local/mysql/data/ibdata*

-rw-rw---- 1 mysql mysql 114M Mar 11 14:00 /usr/local/mysql/data/ibdata1
```

“ ” 114MB InnoDB rollback
InnoDB undo
undo

2 master thread 10 full purge
UPDATE ibdata1

python py_innodb_page_info
code.google.com david-mysql-tools

```
[root@nineyou0-43 py]# python py_innodb_page_info.py /usr/local/mysql/data/ibdata1

Total number of page:83584:

Insert Buffer Free List:204

Freshly Allocated Page:5467

Undo Log Page:38675

File Segment inode:4

B-tree Node:39233

File Space Header:1
```

83 584204546738 675
undo39 233-v

4.2.2 B+

Figure 4-1 illustrates the structure of an InnoDB index organized as a B+ tree. The diagram shows a root node (Non-leaf node segment) pointing to several leaf nodes (Leaf node segment). Each leaf node contains a list of pointers to data blocks. The B+ tree structure is used for efficient index access and range queries.

InnoDB uses a B+ tree structure for its indexes. The root node is a Non-leaf node segment, and the leaf nodes are Leaf node segments. The leaf nodes contain pointers to data blocks. The B+ tree structure is used for efficient index access and range queries.

4.2.3 InnoDB

InnoDB 默认使用 1MB 的 InnoDB 表空间，使用 4 个 5 页的 InnoDB 表空间，16KB 的 InnoDB 表空间，64 个 InnoDB 表空间。

InnoDB 1.0.x 使用 KEY_BLOCK_SIZE 参数，2K、4K、8K、16K、32K、64K、128K、256K、512K、1M。

InnoDB 1.2.x 使用 innodb_page_size 参数，4K、8K、16K、32K、64K、128K、256K、512K、1M。

InnoDB 使用 innodb_file_per_table 参数，96KB、64KB、32KB、16KB、8KB、4KB、2KB、1KB、512B、256B、128B、64B、32B、16B、8B、4B、2B、1B。fragment page 64 个 InnoDB 表空间，undo 表空间，InnoDB 表空间。

```
mysql>CREATE TABLE t1(  
-col1 INT NOT NULL AUTO_INCREMENT,  
-col2 VARCHAR(7000),  
-PRIMARY KEY(col1))ENGINE=InnoDB;  
  
mysql>system ls -lh /usr/local/mysql/data/test/t1.ibd;  
  
-rw-rw---- 1 mysql mysql 96K 10月12 14:59 /usr/local/mysql/data/test/t1.ibd
```

SQL 语句 t1.col2 VARCHAR(7000) 表空间 96KB，SQL 语句 t1.col2 VARCHAR(7000) 表空间 96KB，SQL 语句 t1.col2 VARCHAR(7000) 表空间 96KB。

```
mysql>INSERT t1 SELECT NULL,REPEAT('a',7000);  
  
Query OK, 1 row affected (0.04 sec)
```

Records:1 Duplicates:0 Warnings:0

```
mysql>INSERT into t1 SELECT NULL,REPEAT('a',7000);
```

Query OK,1 row affected(0.01 sec)

Records:1 Duplicates:0 Warnings:0

```
mysql>system ls-lh/usr/local/mysql/data/test/t1.ibd;
```

```
-rw-rw----1 mysql mysql 96K 10 12 16:24/usr/local/mysql/data/test/t1.ibd
```

py_innodb_page_info

```
[root@nineyou0-43 py]#./py_innodb_page_info.py-v/usr/local/mysql/data/test/t1.ibd
```

page offset 00000000,page typeFile Space Header

page offset 00000001,page typeInsert Buffer Bitmap

page offset 00000002,page typeFile Segment inode

page offset 00000003,page typeB-tree Node,page level0000

page offset 00000000,page typeFreshly Allocated Page

page offset 00000000,page typeFreshly Allocated Page

Total number of page:6:

Freshly Allocated Page:2

Insert Buffer Bitmap:1

File Space Header:1

B-tree Node:1

File Segment inode:1

-vpage offset3page
level0

```
mysql>INSERT into t1 SELECT NULL,REPEAT('a',7000);
```

Query OK,1 row affected(0.01 sec)

Records:1 Duplicates:0 Warnings:0

[root@nineyou0-43 py]#./py_innodb_page_info.py-v/usr/local/mysql/data/test/t1.ibd

page offset 00000000,page type[]File Space Header[]

page offset 00000001,page type[]Insert Buffer Bitmap[]

page offset 00000002,page type[]File Segment inode[]

page offset 00000003,page type[]B-tree Node[],page level[]0001[]

page offset 00000004,page type[]B-tree Node[],page level[]0000[]

page offset 00000005,page type[]B-tree Node[],page level[]0000[]

Total number of page:6:

Insert Buffer Bitmap:1

File Space Header:1

B-tree Node:3

File Segment inode:1

page offset[]3page level[]01
B+ B-tree Node[]

60t16332

mysql\DELIMITER//

mysql\CREATE PROCEDURE load_t1(count INT UNSIGNED)

-BEGIN

-DECLARE s INT UNSIGNED DEFAULT 1;

-DECLARE c VARCHAR(7000)DEFAULT REPEAT('a',7000);

-WHILE s<=count DO

-INSERT INTO t1 SELECT NULL,c;

-SET s=s+1;

-END WHILE;

-END;

-//

Query OK,0 rows affected(0.04 sec)

mysql>DELIMITER;

mysql>CALL load_t1(60);

Query OK,1 row affected(1.59 sec)

mysql>SELECT COUNT(*)FROM t1\G;

*****1.r0w*****

count(*):63

1 row in set(0.00 sec)1 row in set(0.00 sec)

mysql>system ls-lh/usr/local/mysql/data/test/t1.ibd;

-rw-rw----1 mysql mysql 576K 10 12 16:56/usr/local/mysql/data/test/t1.ibd

631MB
64py_innodb_page_info.t.ibd

[root@nineyou0-43 py]#./py_innodb_page_info.py-v/usr/local/mysql/data/test/t1.ibd

page offset 00000000,page typeFile Space Header

page offset 00000001,page typeInsert Buffer Bitmap

page offset 00000002,page typeFile Segment inode

page offset 00000003,page typeB-tree Node,page level0001

page offset 00000004,page typeB-tree Node,page level0000

page offset 00000005,page typeB-tree Node,page level0000

page offset 00000006,page typeB-tree Node,page level0000

page offset 00000007,page typeB-tree Node,page level0000

page offset 00000008,page typeB-tree Node,page level0000

page offset 00000009,page typeB-tree Node,page level0000

page offset 0000000a,page typeB-tree Node,page level0000

page offset 0000000b,page typeB-tree Node,page level0000

page offset 0000000c,page typeB-tree Node,page level0000

page offset 0000000d,page typeB-tree Node,page level0000

page offset 0000000e,page type[]B-tree Node[],page level[]0000[]
page offset 0000000f,page type[]B-tree Node[],page level[]0000[]
page offset 00000010,page type[]B-tree Node[],page level[]0000[]
page offset 00000011,page type[]B-tree Node[],page level[]0000[]
page offset 00000012,page type[]B-tree Node[],page level[]0000[]
page offset 00000013,page type[]B-tree Node[],page level[]0000[]
page offset 00000014,page type[]B-tree Node[],page level[]0000[]
page offset 00000015,page type[]B-tree Node[],page level[]0000[]
page offset 00000016,page type[]B-tree Node[],page level[]0000[]
page offset 00000017,page type[]B-tree Node[],page level[]0000[]
page offset 00000018,page type[]B-tree Node[],page level[]0000[]
page offset 00000019,page type[]B-tree Node[],page level[]0000[]
page offset 0000001a,page type[]B-tree Node[],page level[]0000[]
page offset 0000001b,page type[]B-tree Node[],page level[]0000[]
page offset 0000001c,page type[]B-tree Node[],page level[]0000[]
page offset 0000001d,page type[]B-tree Node[],page level[]0000[]
page offset 0000001e,page type[]B-tree Node[],page level[]0000[]
page offset 0000001f,page type[]B-tree Node[],page level[]0000[]
page offset 00000020,page type[]B-tree Node[],page level[]0000[]
page offset 00000021,page type[]B-tree Node[],page level[]0000[]
page offset 00000022,page type[]B-tree Node[],page level[]0000[]
page offset 00000023,page type[]B-tree Node[],page level[]0000[]

Total number of page:36:

Insert Buffer Bitmap:1

File Space Header:1

B-tree Node:33

File Segment inode:1

[][][]B-tree Node[][][]33[][][]page level[]1[][][]32
[]page level[]0[][][]32[][][]

64

```
mysql>CALL load_t1(1);
```

```
Query OK,1 row affected(0.10 sec)
```

```
mysql>system ls-lh/usr/local/mysql/data/test/t1.ibd;
```

```
-rw-rw----1 mysql mysql 2.0M 10 12 17:02/usr/local/mysql/data/test/t1.ibd
```

32 py_innodb_page_info t1.ibd Freshly Allocated Page

```
[root@nineyou0-43 test2]#py/py_innodb_page_info.py t1.ibd-v
```

```
page offset 00000000,page typeFile Space Header
```

```
page offset 00000001,page typeInsert Buffer Bitmap
```

```
page offset 00000002,page typeFile Segment inode
```

```
page offset 00000003,page typeB-tree Node,page level0001
```

```
page offset 00000004,page typeB-tree Node,page level0000
```

```
page offset 00000005,page typeB-tree Node,page level0000
```

```
page offset 00000006,page typeB-tree Node,page level0000
```

```
page offset 00000007,page typeB-tree Node,page level0000
```

```
page offset 00000008,page typeB-tree Node,page level0000
```

```
page offset 00000009,page typeB-tree Node,page level0000
```

```
page offset 0000000a,page typeB-tree Node,page level0000
```

```
page offset 0000000b,page typeB-tree Node,page level0000
```

```
page offset 0000000c,page typeB-tree Node,page level0000
```

```
page offset 0000000d,page typeB-tree Node,page level0000
```

```
page offset 0000000e,page typeB-tree Node,page level0000
```

```
page offset 0000000f,page typeB-tree Node,page level0000
```

```
page offset 00000010,page typeB-tree Node,page level0000
```

```
page offset 00000011,page typeB-tree Node,page level0000
```

page offset 00000012,page type[]B-tree Node[],page level[]0000[]
page offset 00000013,page type[]B-tree Node[],page level[]0000[]
page offset 00000014,page type[]B-tree Node[],page level[]0000[]
page offset 00000015,page type[]B-tree Node[],page level[]0000[]
page offset 00000016,page type[]B-tree Node[],page level[]0000[]
page offset 00000017,page type[]B-tree Node[],page level[]0000[]
page offset 00000018,page type[]B-tree Node[],page level[]0000[]
page offset 00000019,page type[]B-tree Node[],page level[]0000[]
page offset 0000001a,page type[]B-tree Node[],page level[]0000[]
page offset 0000001b,page type[]B-tree Node[],page level[]0000[]
page offset 0000001c,page type[]B-tree Node[],page level[]0000[]
page offset 0000001d,page type[]B-tree Node[],page level[]0000[]
page offset 0000001e,page type[]B-tree Node[],page level[]0000[]
page offset 0000001f,page type[]B-tree Node[],page level[]0000[]
page offset 00000020,page type[]B-tree Node[],page level[]0000[]
page offset 00000021,page type[]B-tree Node[],page level[]0000[]
page offset 00000022,page type[]B-tree Node[],page level[]0000[]
page offset 00000023,page type[]B-tree Node[],page level[]0000[]
page offset 00000000,page type[]Freshly Allocated Page[]

.....

page offset 00000000,page type[]Freshly Allocated Page[]
page offset 00000000,page type[]Freshly Allocated Page[]
page offset 00000000,page type[]Freshly Allocated Page[]

Total number of page:128:

Freshly Allocated Page:91

Insert Buffer Bitmap:1

File Space Header:1

B-tree Node:34

File Segment inode:1

4.2.4 □

```

InnoDBPageInnoDB
InnoDB16KBInnoDB 1.2.x
innodb_page_size4K8K16K
innodb_page_sizemysqldump

```

□InnoDB□□□□□□□□□□□□□□□□

□□□□B-tree Node□

undo undo Log Page

System Page

Transaction system Page

□□□□□□□□Insert Buffer Bitmap□

□□□□□□□□□□Insert Buffer Free List□

Uncompressed BLOB Page

compressed BLOB Page

4.2.5 列存式数据库

InnoDB 是典型的行存式数据库 (row-oriented)，其数据块大小通常为 16KB，而列存式数据库 (column-oriented) 的数据块大小通常为 16KB/2-200 字节。列存式数据库的典型代表有 MySQL 的 infobright、Sybase IQ、Google Big Table 等。列存式数据库的优点是：1. 存储空间利用率高，因为只需要存储数据变化的部分；2. 查询效率高，因为只需要扫描相关的列；3. 压缩率高，因为重复的数据可以存储一次。

4.3 InnoDB

InnoDBはOracle、Microsoft SQL Serverと異なり、
100%のACID特性を持つ。InnoDB 1.0.xは
InnoDBのCompact、Redundantの2つのモードで動作する。
RedundantはInnoDBのデフォルトモードで、
PHYSICAL RECORD、NEW STYLEのPHYSICAL
RECORD、OLD STYLEのMySQL 5.1の
CompactのSHOW TABLE STATUS
LIKE 'table_name' row_format
の出力結果は

```
mysql> SHOW TABLE STATUS like 'mytest%'\G;
```

```
*****1. row*****
```

```
Name: mytest
```

```
Engine: InnoDB
```

```
Version: 10
```

```
Row_format: Compact
```

```
Rows: 6
```

```
Avg_row_length: 2730
```

```
Data_length: 16384
```

```
Max_data_length: 0
```

```
Index_length: 0
```

```
Data_free: 0
```

```
Auto_increment: NULL
```

```
Create_time: 2009-03-17 13:33:50
```

```
Update_time: NULL
```

```
Check_time: NULL
```

```
Collation: latin1_swedish_ci
```

```
Checksum: NULL
```

```
Create_options:
```

Comment:

*****2.row*****

Name:mytest2

Engine:InnoDB

Version:10

Row_format:Redundant

Rows:0

Avg_row_length:0

Data_length:16384

Max_data_length:0

Index_length:0

Data_free:0

Auto_increment:NULL

Create_time:2009-03-17 13:57:23

Update_time:NULL

Check_time:NULL

Collation:latin1_swedish_ci

Checksum:NULL

Create_options:row_format=REDUNDANT

Comment:

2 rows in set(0.00 sec)

mytestCompactmytest2Redundant
py_innodb_page_info

4.3.1 Compact

CompactMySQL 5.0
4-2Compact

| | | | | | |
|----------|---------|-------|------|------|-------|
| 变长字段长度列表 | NULL标志位 | 记录头信息 | 列1数据 | 列2数据 | |
|----------|---------|-------|------|------|-------|

4-2 Compact

4-2 Compact NULL

255 1

255 2

2 MySQL VARCHAR 65535 NULL NULL 1 record header 5 40 4-1

Compact

```
mysql>CREATE TABLE mytest(
- t1 VARCHAR(10),
- t2 VARCHAR(10),
- t3 CHAR(10),
- t4 VARCHAR(10)
- )ENGINE=INNODB CHARSET=LATIN1 ROW_FORMAT=COMPACT;
```

Query OK,0 rows affected(0.00 sec)

```
mysql>INSERT INTO mytest
- VALUES('a','bb','bb','ccc');
```

Query OK,1 row affected(0.01 sec)

```
mysql>INSERT INTO mytest
- VALUES('d','ee','ee','fff');
```

Query OK,1 row affected(0.00 sec)

```
mysql>INSERT INTO mytest
- VALUES('d',NULL,NULL,'fff');
```

Query OK,1 row affected(0.00 sec)

```
mysql>SELECT*FROM mytest\G;
```

*****1.row*****

t1:a

t2:bb

t3:bb

t4:ccc

*****2.row*****

t1:d

t2:ee

t3:ee

t4:fff

*****3.row*****

t1:d
t2:NULL
t3:NULL
t4:fff

3 rows in set(0.00 sec)

mysql> show create table mytest\G
Create Table mytest (t1 VARCHAR(4), t2 VARCHAR(3), t3 CHAR(3), t4 CHAR(3)) ENGINE=InnoDB
innodb_file_per_table=1

Windows 下使用 UltraEdit 打开 mytest.ibd 文件
hexdump -C -v mytest.ibd
mytest.txt

```
0000c070 73 75 70 72 65 6d 75 6d 03 02 01 00 00 00 10 00|supremum.....|
0000c080 2c 00 00 00 2b 68 00 00 00 00 06 05 80 00 00 00|,...+h.....|
0000c090 00 32 01 10 61 62 62 62 20 20 20 20 20 20 20 20|.2...abbbb|
0000c0a0 20 63 63 63 03 02 01 00 00 00 18 00 2b 00 00 00|ccc.....+...|
0000c0b0 2b 68 01 00 00 00 06 06 80 00 00 00 32 01 10 00|+h.....2..|
0000c0c0 64 65 65 65 65 20 20 20 20 20 20 20 20 66 66 66|deeeefff|
0000c0d0 03 01 06 00 00 20 ff 98 00 00 00 2b 68 02 00 00|.....+h...|
0000c0e0 00 00 06 07 80 00 00 00 32 01 10 64 66 66 66 00|.....2..dfff.|
```

0000c078

```
03 02 01/*Record Header*/
00/*NULL*/
00 00 10 00 2c/*Record Header*/
00 00 00 2b 68 00/*RowID InnoDB*/
00 00 00 00 06 05/*TransactionID*/
80 00 0000 32 01 10/*Roll Pointer*/
```


61/*1'a'*/

62 62/*2'bb'*/

62 62 20 20 20 20 20 20 20/*3'bb'*/

63 63 63/*4'ccc'*/

03 02 0101 02 03InnoDBTransactionID
Roll PointerCHAR0x20

Record Headernext_recorder0x2c
0x2cInnoDB

RowID
NULL

03 01/*'*/

06/*NULL'*/

00 00 20 ff 98/*Record Header*/

00 00 00 2b 68 02/*RowID*/

00 00 00 00 06 07/*TransactionID*/

80 00 00 00 32 01 10/*Roll Pointer*/

64/*1'd'*/

66 66 66/*4'fff'*/

NULLNULL0006000001101
23NULLNULL
14NULLCHARVARCHAR
compactNULL

4.3.2 Redundant

Redundant MySQL 5.0 InnoDB MySQL 5.0
Redundant Redundant 4-3

| | | | | | |
|----------|-------|------|------|------|-------|
| 字段长度偏移列表 | 记录头信息 | 列1数据 | 列2数据 | 列3数据 | |
|----------|-------|------|------|------|-------|

4-3 Redundant

4-3 Compact Redundant
255 1 255 2
record header Compact
Redundant 6 48 4-2 4-2
n_fields 10 MySQL
1023 1 byte_offs_flags 1 2

表 4-2 Redundant 记录头信息

| 名 称 | 大小 (bit) | 描 述 |
|-----------------|----------|-----------------------|
| () | 1 | 未知 |
| () | 1 | 未知 |
| deleted_flag | 1 | 该行是否已被删除 |
| min_rec_flag | 1 | 为 1，如果该记录是预先被定义为最小的记录 |
| n_owned | 4 | 该记录拥有的记录数 |
| heap_no | 13 | 索引堆中该条记录的索引号 |
| n_fields | 10 | 记录中列的数量 |
| lbyte_offs_flag | 1 | 偏移列表为 1 字节还是 2 字节 |
| next_record | 16 | 页中下一条记录的相对位置 |
| Total | 48 | |

4.3.1 mytest Redundant mytest2

```
mysql>CREATE TABLE mytest2
-ENGINE=InnoDB ROW_FORMAT=Redundant
```

-AS

-SELECT*FROM mytest;

Query OK,3 rows affected(0.00 sec)

Records:3 Duplicates:0 Warnings:0

mysqlSHOW TABLE STATUS LIKE'mytest2'\G;

*****1.row*****

Name:mytest2

Engine:InnoDB

Version:10

Row_format:Redundant

Rows:3

Avg_row_length:5461

Data_length:16384

Max_data_length:0

Index_length:0

Data_free:0

Auto_increment:NULL

Create_time:2009-03-18 15:49:42

Update_time:NULL

Check_time:NULL

Collation:latin1_swedish_ci

Checksum:NULL

Create_options:row_format=REDUNDANT

Comment:

1 row in set(0.00 sec)

mysqlSELECT*FROM mytest2\G;

*****1.row*****

t1:a

t2:bb

t3:bb

t4:ccc

*****2.row*****

t1:d

t2:ee

t3:ee

t4:fff

*****3.row*****

t1:d

t2:NULL

t3:NULL

t4:fff

3 rows in set(0.00 sec)

row_formatRedundant hexdump
mytest2.ibd mytest2.txt

0000c070 08 03 00 00 73 75 70 72 65 6d 75 6d 00 23 20 16|....supremum.#.|
0000c080 14 13 0c 06 00 00 10 0f 00 ba 00 00 00 2b 68 0b|.....+h.|
0000c090 00 00 00 00 06 53 80 00 00 00 32 01 10 61 62 62|....S....2..abb|
0000c0a0 62 62 20 20 20 20 20 20 20 63 63 63 23 20 16|bb ccc#.|
0000c0b0 14 13 0c 06 00 00 18 0f 00 ea 00 00 00 2b 68 0c|.....+h.|
0000c0c0 00 00 00 00 06 53 80 00 00 00 32 01 1e 64 65 65|....S....2..dee|
0000c0d0 65 65 20 20 20 20 20 20 20 66 66 66 21 9e 94|ee fff!...|
0000c0e0 14 13 0c 06 00 00 20 0f 00 74 00 00 00 2b 68 0d|.....t...+h.|
0000c0f0 00 00 00 00 06 53 80 00 00 00 32 01 2c 64 00 00|....S....2.,d..|
0000c100 00 00 00 00 00 00 00 66 66 66 00 00 00 00 00|.....fff.....|

23 20 16 14 13 0c 06/**/

00 00 10 0f 00 ba/*Record Header0006000*/

00 00 00 2b 68 0b/*RowID*/

00 00 00 00 06 53/*TransactionID*/

80 00 00 00 32 01 10/*Roll Point*/

61/*0100'a'*/

62 62/*0200'bb'*/

62 62 20 20 20 20 20 20 20/*0300'bb'Char000*/

63 63 63/*0400'ccc'*/

23 20 16 14 13 0c 0600000600c0130140160200230000000000
060000000606+6=0x0C000000000706+6+7=0x13000000001
06+6+7+1=0x1400000000206+6+7+1+2=0x160000000010
06+6+7+1+2+10=0x20000000003
06+6+7+1+2+10+3=0x2300

000000000000Record Header00000048000022032000
0000000111000000700000000030000000003300100000000000
000

0000000000000000000000Redundant0000000000000000000000
000000000000NULL00000000000000

21 9e 94 14 13 0c 06/*000000000000*/

00 00 20 0f 00 74/*Record Header0006000*/

00 00 00 2b 68 0d/*RowID*/

00 00 00 00 06 53/*TransactionID*/

80 00 00 00 32 01 10/*Roll Point*/

64/*0100'd'*/

00 00 00 00 00 00 00 00 00/*0300NULL*/

66 66 66/*0400'fff'*/

Compact 06 0c
13 14 94 9e 21 4 5 NULL 94 6 CHAR
NULL 9e 94+10=0x9e 21 14+3=0x21
VARCHAR NULL Redundant CHAR
NULL

mytest2 Latin1 1 mytest2
utf8 CHAR 10 10×3=30
Redundant CHAR

4.3.3 五五五五五

```
InnoDB BLOB LOB
BLOB
VARCHAR
```

```

VARCHAR(4000)DBA MySQL VARCHAR(4000)
Oracle VARCHAR2(4000)SQL Server(8000)
MySQL VARCHAR(65535)
65535 VARCHAR(65535)

```

```
mysql>CREATE TABLE test(
```

- a VARCHAR(65535)

```
-) CHARSET=latin1 ENGINE=InnoDB;
```

```
ERROR 1118(42000):Row size too large.The maximum row size for the used table type,not counting BLOBs,is 65535.You have to change some columns to TEXT or BLOBs
```

```

InnoDB[65535]VARCHAR[65532]

```

```
mysql> CREATE TABLE test(
```

- a VARCHAR(65532)

```
-) CHARSET=latin1 ENGINE=InnoDB;
```

Query OK, 0 rows affected (0.15 sec)

```

SQL_MODE=
MySQL warning

```

```
mysql> CREATE TABLE test(
```

- a VARCHAR(65535)

```
-) CHARSET=latin1 ENGINE=InnoDB;
```


Query OK, 0 rows affected, 1 warning (0.14 sec)

mysql> SHOW WARNINGS\G;

*****1. row*****

Level: Note

Code: 1246

Message: Converting column 'a' from VARCHAR to TEXT

1 row in set (0.00 sec)

warning MySQL VARCHAR TEXT test

mysql> SHOW CREATE TABLE test\G;

*****1. row*****

Table: test

Create Table: CREATE TABLE 'test' (

'a' mediumtext

) ENGINE=InnoDB DEFAULT CHARSET=utf8

1 row in set (0.00 sec)

VARCHAR 65 532 latin1 GBK UTF-8

mysql> CREATE TABLE test(

-a VARCHAR(65532)

-) CHARSET=GBK ENGINE=InnoDB;

ERROR 1074 (42000): Column length too big for column 'a' (max=32767); use BLOB or TEXT instead

mysql> mysql> CREATE TABLE test(

-a VARCHAR(65532)

-) CHARSET=UTF8 ENGINE=InnoDB;

ERROR 1074 (42000): Column length too big for column 'a' (max=21845); use BLOB or TEXT instead

MySQL 5.6.21 版本中，VARCHAR 类型的最大长度是 65535 字节，而不是 65532 字节。这是因为 VARCHAR 类型的长度是动态的，它会根据存储的字符串长度而变化。在 MySQL 5.6.21 之前，VARCHAR 类型的最大长度是 65532 字节，这是因为当时的实现中，VARCHAR 类型的长度是固定的，它会根据存储的字符串长度而变化。在 MySQL 5.6.21 之后，VARCHAR 类型的最大长度是 65535 字节，这是因为当时的实现中，VARCHAR 类型的长度是动态的，它会根据存储的字符串长度而变化。

MySQL 5.6.21 版本中，VARCHAR 类型的最大长度是 65535 字节，而不是 65532 字节。这是因为 VARCHAR 类型的长度是动态的，它会根据存储的字符串长度而变化。在 MySQL 5.6.21 之前，VARCHAR 类型的最大长度是 65532 字节，这是因为当时的实现中，VARCHAR 类型的长度是固定的，它会根据存储的字符串长度而变化。在 MySQL 5.6.21 之后，VARCHAR 类型的最大长度是 65535 字节，这是因为当时的实现中，VARCHAR 类型的长度是动态的，它会根据存储的字符串长度而变化。

```
mysql> CREATE TABLE test2(
```

```
- a VARCHAR(22000),
```

```
- b VARCHAR(22000),
```

```
- c VARCHAR(22000)
```

```
); CHARSET=latin1 ENGINE=InnoDB;
```

ERROR 1118 (42000): Row size too large. The maximum row size for the used table type, not counting BLOBs, is 65535. You have to change some columns to TEXT or BLOBs

MySQL 5.6.21 版本中，InnoDB 引擎的 VARCHAR 类型的最大长度是 65532 字节，而不是 65535 字节。这是因为 InnoDB 引擎的 VARCHAR 类型的长度是固定的，它会根据存储的字符串长度而变化。在 MySQL 5.6.21 之前，InnoDB 引擎的 VARCHAR 类型的最大长度是 65532 字节，这是因为当时的实现中，InnoDB 引擎的 VARCHAR 类型的长度是固定的，它会根据存储的字符串长度而变化。在 MySQL 5.6.21 之后，InnoDB 引擎的 VARCHAR 类型的最大长度是 65535 字节，这是因为当时的实现中，InnoDB 引擎的 VARCHAR 类型的长度是动态的，它会根据存储的字符串长度而变化。

```
mysql> CREATE TABLE t(
```

```
- a VARCHAR(65532)
```

```
); ENGINE=InnoDB CHARSET=latin1;
```

Query OK, 0 rows affected (0.15 sec)

```
mysql> INSERT INTO t
```

```
- SELECT REPEAT('a', 65532);
```

Query OK, 1 row affected (0.08 sec)

Records: 1 Duplicates: 0 Warnings: 0

MySQL 5.6.21 版本中，InnoDB 引擎的 VARCHAR 类型的最大长度是 65532 字节，而不是 65535 字节。这是因为 InnoDB 引擎的 VARCHAR 类型的长度是固定的，它会根据存储的字符串长度而变化。在 MySQL 5.6.21 之前，InnoDB 引擎的 VARCHAR 类型的最大长度是 65532 字节，这是因为当时的实现中，InnoDB 引擎的 VARCHAR 类型的长度是固定的，它会根据存储的字符串长度而变化。在 MySQL 5.6.21 之后，InnoDB 引擎的 VARCHAR 类型的最大长度是 65535 字节，这是因为当时的实现中，InnoDB 引擎的 VARCHAR 类型的长度是动态的，它会根据存储的字符串长度而变化。

page offset 00000001,page type[]Insert Buffer Bitmap[]

```
page offset 00000003,page typeB-tree Node,page level0000
```

page offset 00000005,page type[Uncompressed BLOB Page]

page offset 00000007,page type Uncompressed BLOB Page

Insert Buffer Bitmap:1

File Space Header:1

File Segment inode:1

```
0000c010 00 00 00 0a 6a d9 c0 89 45 bf 00 00 00 00 00 00|...j...E.....|
```

```
0000c030 00 80 00 05 00 00 00 01 00 00 00 00 00 00 00|.....|
```

```
0000c050 00 02 00 f2 00 00 00 c3 00 00 00 02 00 32 01 00|.....2..|
```

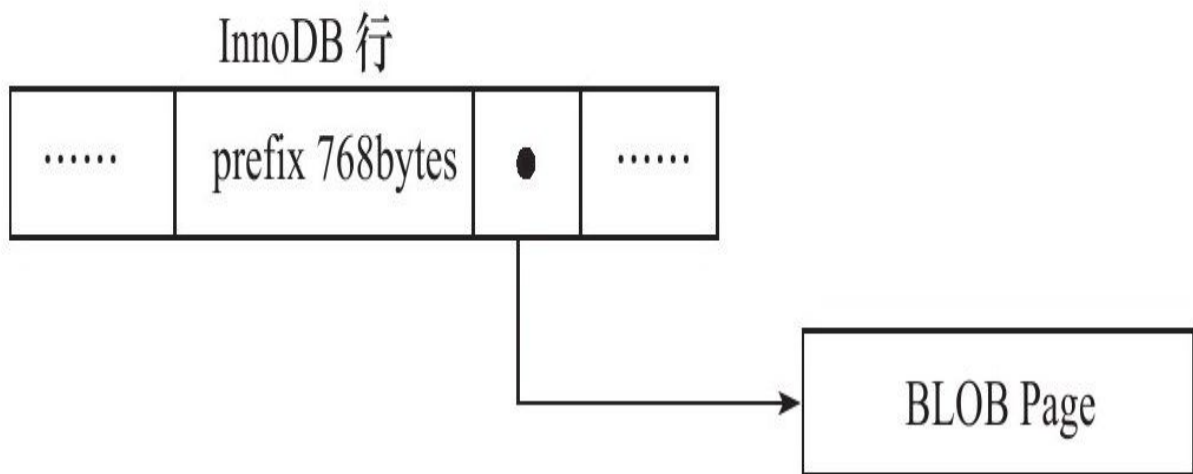
```
0000c070 73 75 70 72 65 6d 75 6d 14 c3 00 00 00 10 ff f0|supremum.....|
```

```

0000c090 2d 01 10 61 61 61 61 61 61 61 61 61 61 61|-...aaaaaaaaaaaa|
0000c0a0 61 61 61 61 61 61 61 61 61 61 61 61 61 61|aaaaaaaaaaaaaaaa|
0000c0b0 61 61 61 61 61 61 61 61 61 61 61 61 61 61|aaaaaaaaaaaaaaaa|
0000c0c0 61 61 61 61 61 61 61 61 61 61 61 61 61 61|aaaaaaaaaaaaaaaa|
0000c0d0 61 61 61 61 61 61 61 61 61 61 61 61 61 61|aaaaaaaaaaaaaaaa|
0000c0e0 61 61 61 61 61 61 61 61 61 61 61 61 61 61|aaaaaaaaaaaaaaaa|
0000c0f0 61 61 61 61 61 61 61 61 61 61 61 61 61 61|aaaaaaaaaaaaaaaa|
0000c100 61 61 61 61 61 61 61 61 61 61 61 61 61 61|aaaaaaaaaaaaaaaa|
0000c110 61 61 61 61 61 61 61 61 61 61 61 61 61 61|aaaaaaaaaaaaaaaa|
.....
0000c390 61 61 61 00 00 00 c3 00 00 00 04 00 00 00 26 00|aaa.....|.
0000c3a0 00 00 00 00 00 fc fc 00 00 00 00 00 00 00 00|.....|

```

0x0000c093 0x0000c392 VARCHA
 65532 768 prefix a
 Uncompressed BLOB Page 4-4
 4-4



4-4

VARCHAR BLOB InnoDB B+Tree B+Tree InnoDB

```
mysql>CREATE TABLE t(
-a VARCHAR(9000)
-)ENGINE=InnoDB;

Query OK,0 rows affected(0.13 sec)

mysql>INSERT INTO t
-SELECT REPEAT('a',9000);

Query OK,1 row affected(0.04 sec)

Records:1 Duplicates:0 Warnings:0

mysql>INSERT INTO t
-SELECT REPEAT('a',9000);

Query OK,1 row affected(0.04 sec)

Records:1 Duplicates:0 Warnings:0
```

t a 9000 py_innodb_page_info BLOB

```
[root@nineyou0-43 mytest]#py_innodb_page_info.py-v t.ibd

page offset 00000000,page typeFile Space Header
page offset 00000001,page typeInsert Buffer Bitmap
page offset 00000002,page typeFile Segment inode
page offset 00000003,page typeB-tree Node,page level0000
page offset 00000004,page typeUncompressed BLOB Page
page offset 00000005,page typeUncompressed BLOB Page

Total number of page:6:

Insert Buffer Bitmap:1
```

Uncompressed BLOB Page:2

File Space Header:1

B-tree Node:1

File Segment inode:1

py_innodb_page_info

VARCHAR BLOB
8098 varchar 8098
2

mysql CREATE TABLE t(

-a varchar(8098)

)ENGINE=InnoDB;

Query OK,0 rows affected(0.12 sec)

mysql INSERT INTO t SELECT REPEAT('a',8098);

Query OK,1 row affected(0.04 sec)

Records:1 Duplicates:0 Warnings:0

mysql INSERT INTO t SELECT REPEAT('a',8098);

Query OK,1 row affected(0.03 sec)

Records:1 Duplicates:0 Warnings:0

py_innodb_page_info t.ibd
BLOB Microsoft SQL Server DBA
InnoDB VARCHAR SQL Server varchar MAX

[root@nineyou0-43 mytest]#py_innodb_page_info.py-v t.ibd

page offset 00000000,page typeFile Space Header

page offset 00000001,page typeInsert Buffer Bitmap

page offset 00000002,page type[]File Segment inode[]
page offset 00000003,page type[]B-tree Node[],page level[]0000[]
page offset 00000000,page type[]Freshly Allocated Page[]
page offset 00000000,page type[]Freshly Allocated Page[]
Total number of page:6:
Freshly Allocated Page:2
Insert Buffer Bitmap:1
File Space Header:1
B-tree Node:1
File Segment inode:1

TEXT BLOB Uncompressed
BLOB Page BLOB
VARCHAR

```
mysql>CREATE TABLE t(
-a BLOB
-)ENGINE=InnoDB;

Query OK,0 rows affected(0.12 sec)

mysql>INSERT INTO t SELECT REPEAT('a',8000);

Query OK,1 row affected(0.03 sec)

Records:1 Duplicates:0 Warnings:0

mysql>INSERT INTO t SELECT REPEAT('a',8000);

Query OK,1 row affected(0.03 sec)

Records:1 Duplicates:0 Warnings:0

mysql>INSERT INTO t SELECT REPEAT('a',8000);

Query OK,1 row affected(0.01 sec)

Records:1 Duplicates:0 Warnings:0

mysql>INSERT INTO t SELECT REPEAT('a',8000);

Query OK,1 row affected(0.06 sec)
```

Records:1 Duplicates:0 Warnings:0

py_innodb_page_info.py 4 8000
py_innodb_page_info.py t.ibd BLOB

```
[root@nineyou0-43 mytest]#py_innodb_page_info.py -v t.ibd

page offset 00000000,page typeFile Space Header

page offset 00000001,page typeInsert Buffer Bitmap

page offset 00000002,page typeFile Segment inode

page offset 00000003,page typeB-tree Node,page level0001

page offset 00000004,page typeB-tree Node,page level0000

page offset 00000005,page typeB-tree Node,page level0000

page offset 00000006,page typeB-tree Node,page level0000

page offset 00000000,page typeFreshly Allocated Page

Total number of page:8:

Freshly Allocated Page:1

Insert Buffer Bitmap:1

File Space Header:1

B-tree Node:4

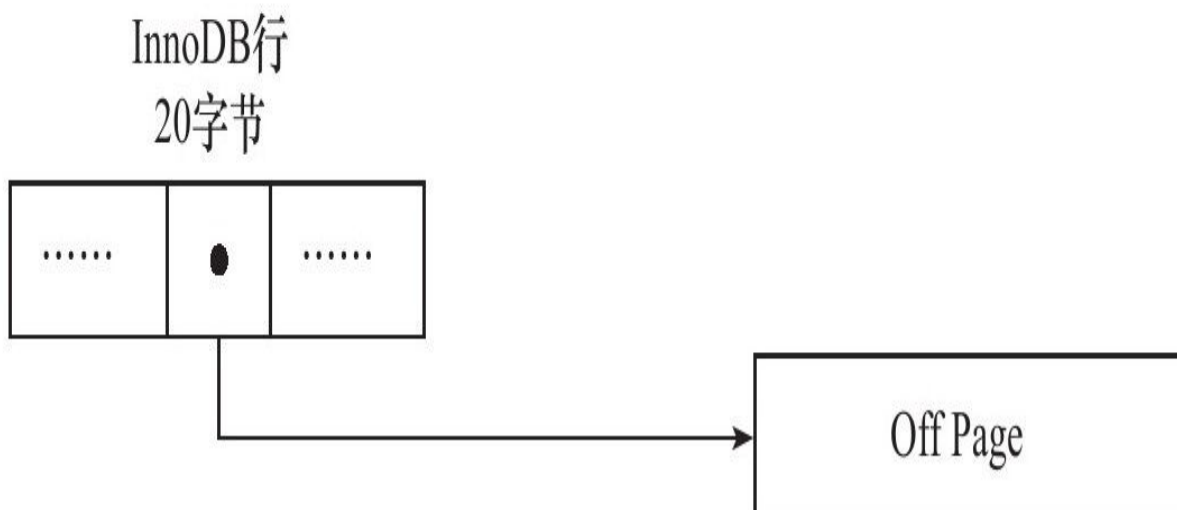
File Segment inode:1
```

BLOB BLOB
BLOB 768

4.3.4 Compressed Dynamic

InnoDB 1.0.x file format
Compact Redundant Antelope
Barracuda Barracuda Compressed
Dynamic

BLOB 4-5
20 Off Page Compact Redundant
768



4-5 Barracuda

Compressed
BLOB TEXT VARCHAR
zlib

4.3.5 CHAR数据类型

MySQL 4.1 版本之前，CHAR 数据类型只能存储 ASCII 字符。MySQL 4.1 版本之后，CHAR 数据类型可以存储 Latin1 字符。

MySQL 4.1 版本之后，CHAR 数据类型可以存储 Latin1 字符。MySQL 4.1 版本之前，CHAR 数据类型只能存储 ASCII 字符。

```
mysql> CREATE TABLE j(
  -> a CHAR(2)
  ->)CHARSET=GBK ENGINE=InnoDB;
Query OK, 0 rows affected (0.11 sec)

mysql> INSERT INTO j SELECT 'ab';
Query OK, 1 row affected (0.03 sec)

Records: 1 Duplicates: 0 Warnings: 0

mysql> SET NAMES GBK;
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO j SELECT '';
Query OK, 1 row affected (0.04 sec)

Records: 1 Duplicates: 0 Warnings: 0

mysql> INSERT INTO j SELECT 'a';
Query OK, 1 row affected (0.03 sec)

Records: 1 Duplicates: 0 Warnings: 0
```

MySQL 4.1 版本之后，CHAR 数据类型可以存储 Latin1 字符。MySQL 4.1 版本之前，CHAR 数据类型只能存储 ASCII 字符。

```
mysql> SELECT a, CHAR_LENGTH(a), LENGTH(a)
  -> FROM j\G;
```

```
*****1.row*****
```

```
a:ab
```

```
CHAR_LENGTH(a):2
```

```
LENGTH(a):2
```

```
*****2.row*****
```

```
a:
```

```
CHAR_LENGTH(a):2
```

```
LENGTH(a):4
```

```
*****3.row*****
```

```
a:a
```

```
CHAR_LENGTH(a):1
```

```
LENGTH(a):1
```

```
3 rows in set(0.00 sec)
```

```
CHAR_LENGTH CHAR 'ab' ' '
2 'ab' 2 ' ' 4 HEX

```

```
mysql> SELECT a, HEX(a)
```

```
-> FROM j\G;
```

```
*****1.row*****
```

```
a:ab
```

```
HEX(a):6162
```

```
*****2.row*****
```

```
a:
```

```
HEX(a):CED2C3C7
```

```
*****3.row*****
```

```
a:a
```

```
HEX(a):61
```

```
3 rows in set(0.00 sec)
```

0xab'0x6162'0xCED2C3C7
CHARUTF-8CHAR10
1030
CHARInnoDB
CHARhexdumpj.ibd

0000c070 73 75 70 72 65 6d 75 6d 02 00 00 00 10 00 1c 00|supremum.....|
0000c080 00 00 b6 2b 2b 00 00 00 51 52 da 80 00 00 00 2d|...++...QR.....-|
0000c090 01 10 61 62 04 00 00 00 18 ff d5 00 00 00 b6 2b|..ab.....+|
0000c0a0 2c 00 00 00 51 52 db 80 00 00 00 2d 01 10 ce d2|,...QR.....-....|
0000c0b0 c3 c7 00 00 00 00 00 00 00 00 00 00 00 00|.....|

02/*2CHAR*/
00/*NULL*/
00 00 10 00 1c/*Recoder Header*/
00 00 00 b6 2b 2b/*RowID*/
00 00 00 51 52 da/*TransactionID*/
80 00 00 00 2d 01 10/*Roll Point*/
61 62/*'ab'*/

04/*4CHAR*/
00/*NULL*/
00 00 18 ff d5/*Recoder Header*/
00 00 00 b6 2b 2c/*RowID*/
00 00 00 51 52 db/*TransactionID*/
80 00 00 00 2d 01 10/*Roll Point*/

c3 d2 c3 c7/*'*/

#

02/*2CHAR*/

00/*NULL*/

00 00 20 ff b7/*Recoder Header*/

00 00 00 b6 2b 2d/*RowID*/

00 00 00 51 53 17/*TransactionID*/

80 00 00 00 2d 01 10/*Roll Point*/

61 20/*'a'*/

InnoDBCHARCHAR
0x20InnoDB
HEXCHARVARCHAR

4.4 InnoDB

InnoDB 使用 B-tree Node 来存储数据。InnoDB 的存储结构如下：

InnoDB 的存储结构如下：MySQL 的 InnoDB 存储引擎使用 Peter 的 Compact 存储结构。

InnoDB 的存储结构如下：

File Header

Page Header

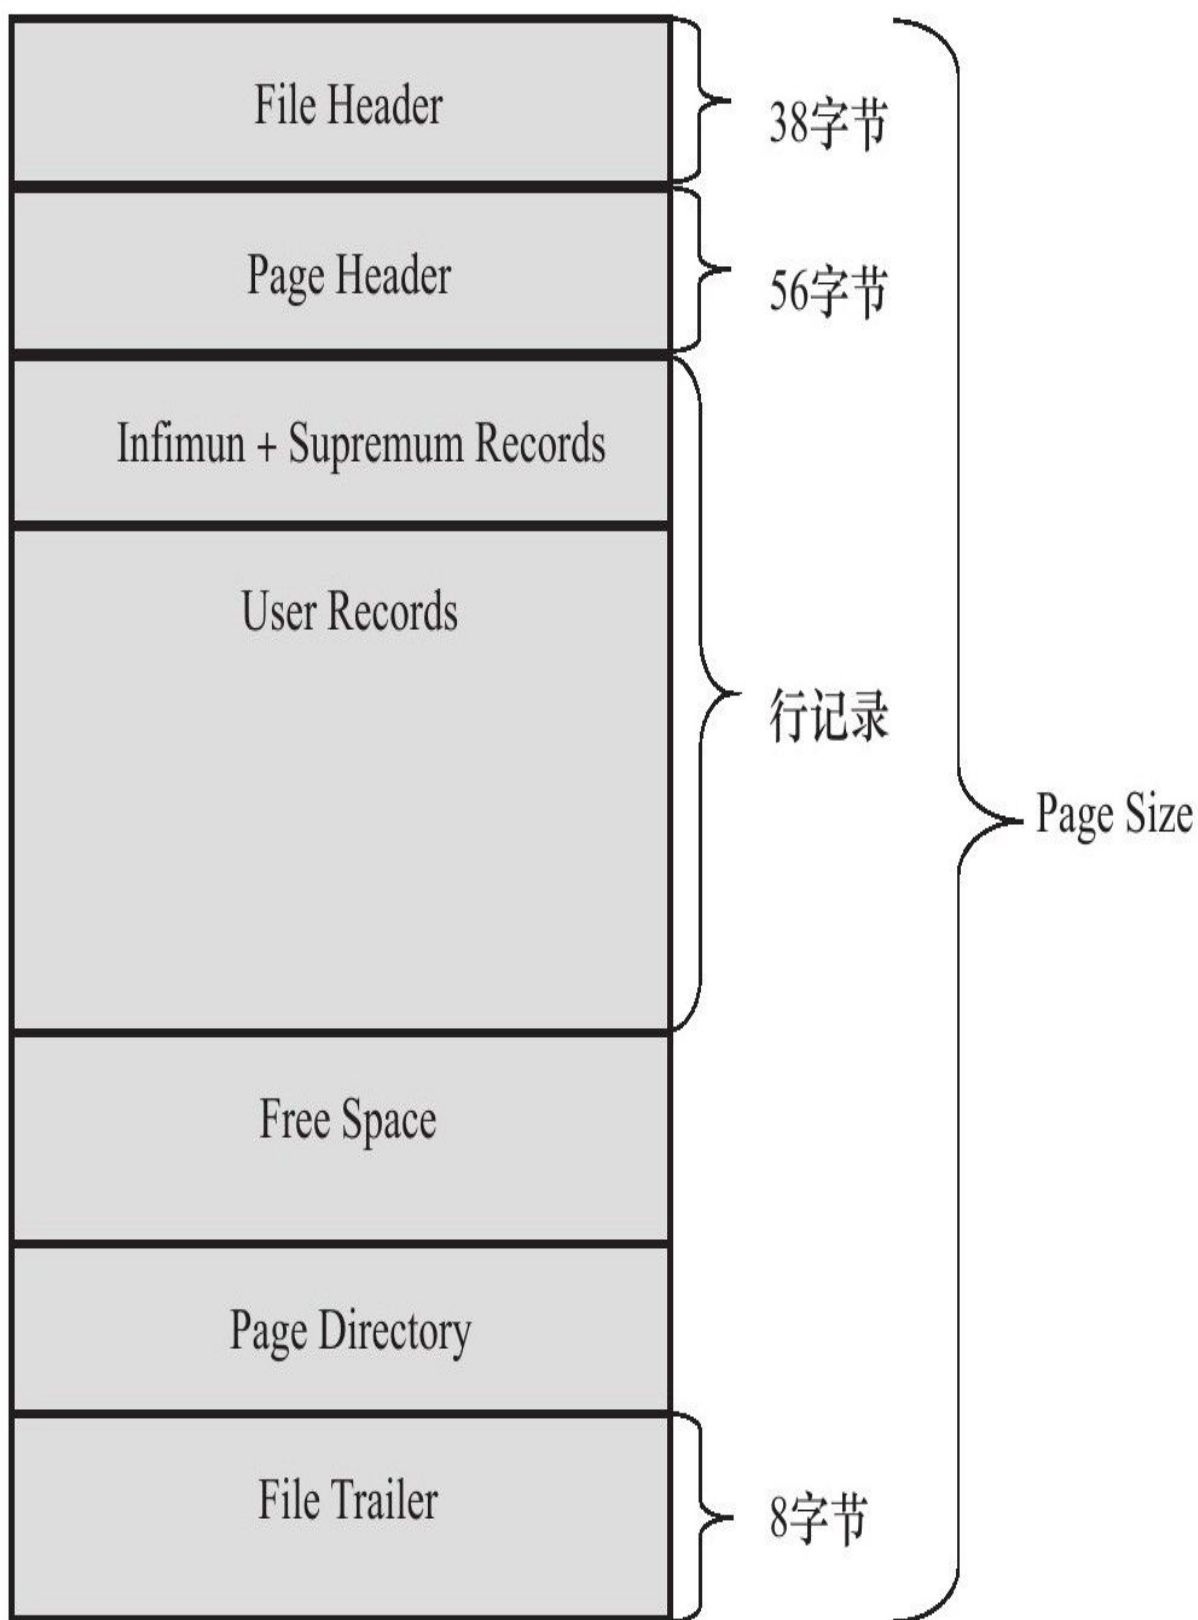
Infimum/Supremum Records

User Records

Free Space

Page Directory

File Trailer



4-6 InnoDB

File HeaderPage HeaderFile TrailerChecksumB+User
RecordsFree SpacePage Directory

4.4.1 File Header

File Header4-38

表 4-3 File Header 组成部分

| 名 称 | 大小 (字节) | 说 明 |
|----------------------------------|---------|---|
| FIL_PAGE_SPACE_OR_CHKSUM | 4 | 当 MySQL 为 MySQL4.0.14 之前的版本时, 该值为 0。在之后的 MySQL 版本中, 该值代表页的 checksum 值 (一种新的 checksum 值) |
| FIL_PAGE_OFFSET | 4 | 表空间中页的偏移值。如某独立表空间 a.ibd 的大小为 1GB, 如果页的大小为 16KB, 那么总共有 65536 个页。FIL_PAGE_OFFSET 表示该页在所有页中的位置。若此表空间的 ID 为 10, 那么搜索页 (10, 1) 就表示查找表 a 中的第二个页 |
| FIL_PAGE_PREV | 4 | 当前页的上一个页, B+ Tree 特性决定了叶子节点必须是双向列表 |
| FIL_PAGE_NEXT | 4 | 当前页的下一个页, B+ Tree 特性决定了叶子节点必须是双向列表 |
| FIL_PAGE_LSN | 8 | 该值代表该页最后被修改的日志序列位置 LSN (Log Sequence Number) |
| FIL_PAGE_TYPE | 2 | InnoDB 存储引擎页的类型。常见的类型见表 4-4。记住 0x45BF, 该值代表了存放的是数据页, 即实际行记录的存储空间 |
| FIL_PAGE_FILE_FLUSH_LSN | 8 | 该值仅在系统表空间的一个页中定义, 代表文件至少被更新到了该 LSN 值。对于独立表空间, 该值都为 0 |
| FIL_PAGE_ARCH_LOG_NO_OR_SPACE_ID | 4 | 从 MySQL 4.1 开始, 该值代表页属于哪个表空间 |

表 4-4 InnoDB 存储引擎中页的类型

| 名 称 | 十六进制 | 解 释 |
|-------------------------|--------|--------------------|
| FIL_PAGE_INDEX | 0x45BF | B+ 树叶节点 |
| FIL_PAGE_UNDO_LOG | 0x0002 | Undo Log 页 |
| FIL_PAGE_INODE | 0x0003 | 索引节点 |
| FIL_PAGE_IBUF_FREE_LIST | 0x0004 | Insert Buffer 空闲列表 |
| FIL_PAGE_TYPE_ALLOCATED | 0x0000 | 该页为最新分配 |
| FIL_PAGE_IBUF_BITMAP | 0x0005 | Insert Buffer 位图 |
| FIL_PAGE_TYPE_SYS | 0x0006 | 系统页 |
| FIL_PAGE_TYPE_TRX_SYS | 0x0007 | 事务系统数据 |
| FIL_PAGE_TYPE_FSP_HDR | 0x0008 | File Space Header |
| FIL_PAGE_TYPE_XDES | 0x0009 | 扩展描述页 |
| FIL_PAGE_TYPE_BLOB | 0x000A | BLOB 页 |

4.4.2 Page Header

File HeaderPage Header14
564-5

表 4-5 Page Header 组成部分

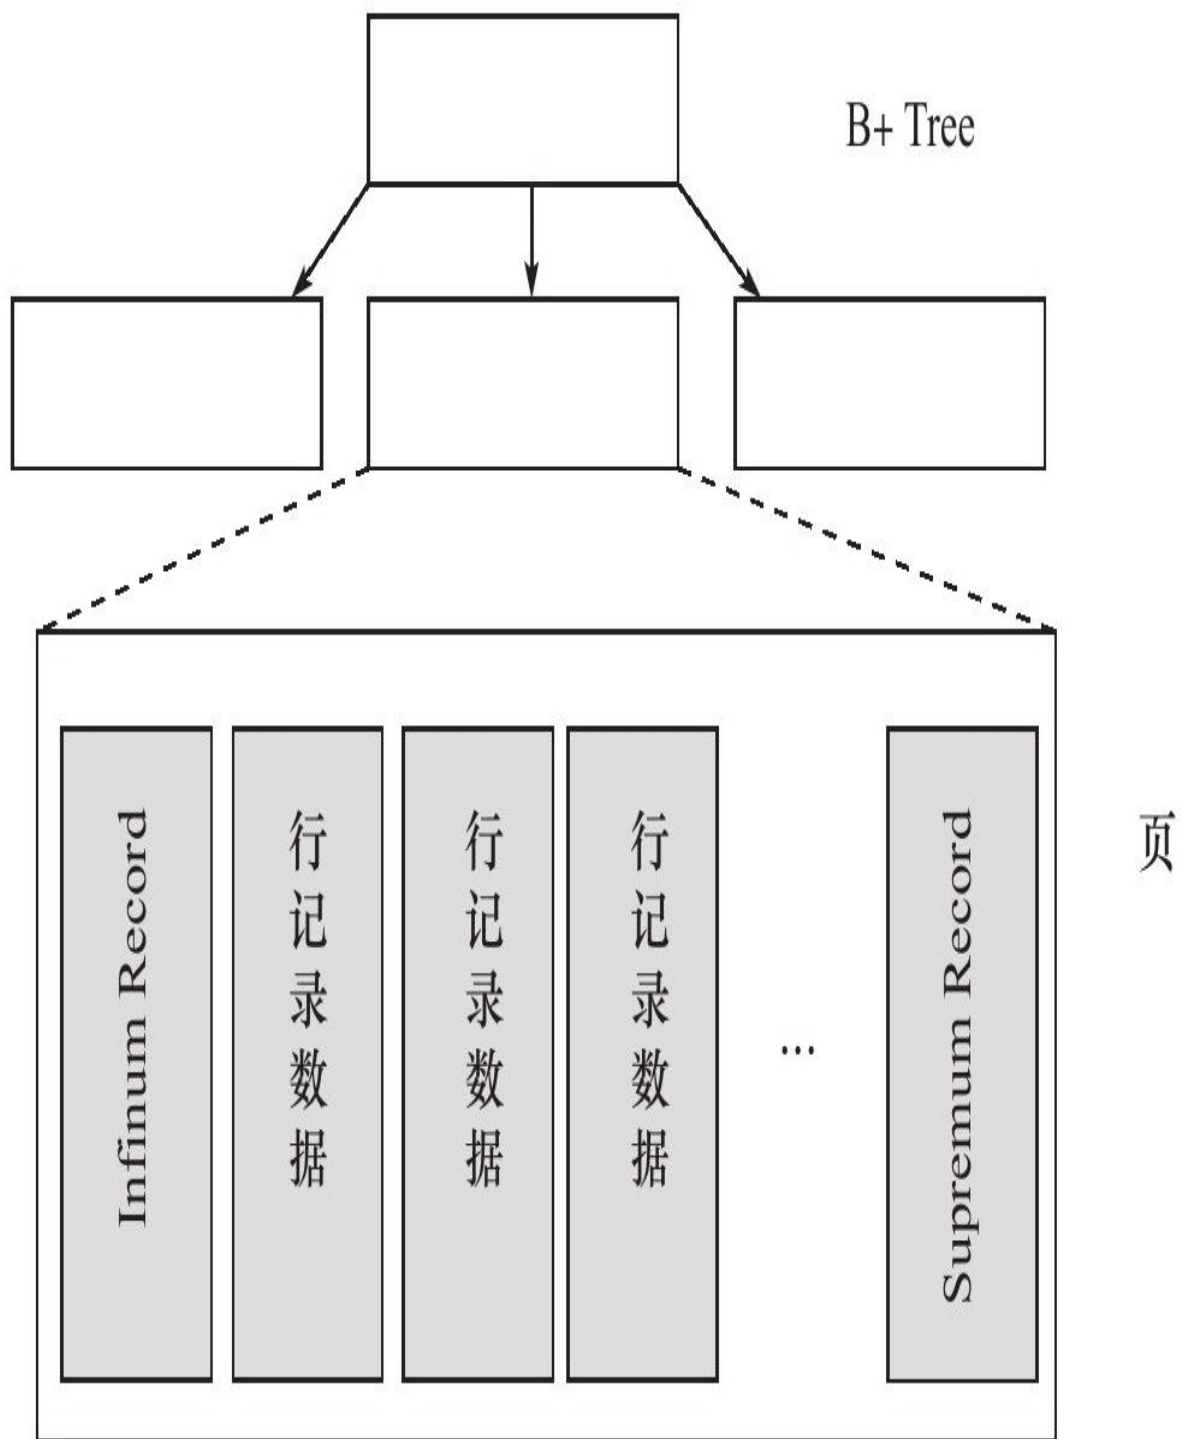
| 名 称 | 大小 (字节) | 说 明 |
|------------------|---------|--|
| PAGE_N_DIR_SLOTS | 2 | 在 Page Directory (页 目 录) 中的 Slot (槽) 数, “4.4.5 Page Directory” 小节中会介绍 |
| PAGE_HEAP_TOP | 2 | 堆中第一个记录的指针, 记录在页中是根据堆的形式存放的 |
| PAGE_N_HEAP | 2 | 堆中的记录数。一共占用 2 字节, 但是第 15 位表示行记录格式 |
| PAGE_FREE | 2 | 指向可重用空间的首指针 |
| PAGE_GARBAGE | 2 | 已删除记录的字节数, 即行记录结构中 delete flag 为 1 的记录大小的总数 |
| PAGE_LAST_INSERT | 2 | 最后插入记录的位置 |
| PAGE_DIRECTION | 2 | 最后插入的方向。可能的取值为:
<input type="checkbox"/> PAGE_LEFT (0x01)
<input type="checkbox"/> PAGE_RIGHT (0x02)
<input type="checkbox"/> PAGE_SAME_REC (0x03)
<input type="checkbox"/> PAGE_SAME_PAGE (0x04)
<input type="checkbox"/> PAGE_NO_DIRECTION (0x05) |
| PAGE_N_DIRECTION | 2 | 一个方向连续插入记录的数量 |
| PAGE_N_RECS | 2 | 该页中记录的数量 |
| PAGE_MAX_TRX_ID | 8 | 修改当前页的最大事务 ID, 注意该值仅在 Secondary Index 中定义 |
| PAGE_LEVEL | 2 | 当前页在索引树中的位置, 0x00 代表叶节点, 即叶节点总是在第 0 层 |
| PAGE_INDEX_ID | 8 | 索引 ID, 表示当前页属于哪个索引 |

(续)

| 名 称 | 大小 (字节) | 说 明 |
|-------------------|---------|---|
| PAGE_BTR_SEG_LEAF | 10 | B+ 树数据页非叶节点所在段的 segment header。注意该值仅在 B+ 树的 Root 页中定义 |
| PAGE_BTR_SEG_TOP | 10 | B+ 树数据页所在段的 segment header。注意该值仅在 B+ 树的 Root 页中定义 |

4.4.3 Infimum-Supremum Record

InnoDB maintains a record of the infimum and supremum of each index. This record is used to determine the range of values that are stored in the index. The record is updated whenever a new value is inserted into the index. The record is also used to determine the range of values that are deleted from the index. The record is updated whenever a value is deleted from the index. The record is also used to determine the range of values that are compacted. The record is updated whenever a value is compacted. The record is also used to determine the range of values that are redundant. The record is updated whenever a value is redundant. The record is also used to determine the range of values that are 4-7. The record is updated whenever a value is 4-7. The record is also used to determine the range of values that are Infimum-Supremum.



□ 4-7 Infinum□Supremum Record

4.4.4 User Record Free Space

User Record InnoDB B+

Free Space

4.4.5 Page Directory

```

Page Directory
Slots
Directory Slots
InnoDB
InnoDB
sparse
directory
Infimum
n_owned
1
Supremum
n_owned
[1,8]
n_owned
[4,8]

```

```

Slots
'i'd'c'b'e'g'l'h'f'j'k'a'4
Slots'a'e'i'

```

```

InnoDBPage Directory
InnoDBrecorder headernext_record
Page Directoryrecorder headern_owned
Page Directory

```

B+ Page Directory

4.4.6 File Trailer

MySQL 5.6.6 之前，InnoDB 的 File Trailer 是 8 字节的。

File Trailer 包含 FIL_PAGE_END_LSN 和 4 字节的 checksum。File Header 包含 FIL_PAGE_LSN 和 4 字节的 checksum。InnoDB 的 checksum 是 4 字节的。如果 checksum 不为 0，则表示文件损坏。not corrupted。

InnoDB 的 Corrupt 标志位。File Trailer 包含 innodb_checksums 标志位。

MySQL 5.6.6 之前，innodb_checksum_algorithm 标志位。checksum 标志位。crc32 标志位。innodb 标志位。strict_innodb 标志位。strict_crc32 标志位。strict_none 标志位。

innodb 标志位。InnoDB 标志位。checksum 标志位。crc32 标志位。MySQL 5.6.6 之前，checksum 标志位。innodb 标志位。checksum 标志位。strict 标志位。MySQL 标志位。none 标志位。checksum 标志位。

strict_* 标志位。checksum 标志位。MySQL 标志位。MySQL 5.6.6 之前，strict_crc32 标志位。innodb 标志位。crc32 标志位。mysql_upgrade 标志位。

4.4.7 InnoDB테이블 생성

테이블 생성을 위해서는 InnoDB 엔진을 지정해야 합니다. 이 예제에서는 테이블 t를 생성합니다.

```
mysql> DROP TABLE IF EXISTS t;

Query OK, 0 rows affected (0.04 sec)

mysql> CREATE TABLE t(
  -a INT UNSIGNED NOT NULL AUTO_INCREMENT,
  -b CHAR(10),
  -PRIMARY KEY(a),
  -)ENGINE=InnoDB CHARSET=utf8;

Query OK, 0 rows affected (0.00 sec)

mysql> DELIMITER $$

mysql> CREATE PROCEDURE load_t(count INT UNSIGNED)
  -BEGIN
  -SET @c=0;
  -WHILE @c<count DO
  -INSERT INTO t
  -SELECT NULL, REPEAT(CHAR(97+RAND()*26), 10);
  -SET @c=@c+1;
  -END WHILE;
  -END;
  -$$

Query OK, 0 rows affected (0.00 sec)

mysql> DELIMITER;

mysql> CALL load_t(100);

Query OK, 0 rows affected (0.60 sec)
```

```
mysql>SELECTa,bFROM t LIMIT 10;
```

```
+----+-----+
```

```
|a|b|
```

```
+----+-----+
```

```
|1|ddddddddd|
```

```
|2|hhhhhhhhh|
```

```
|3|bbbbbbbbb|
```

```
|4|iiiiiii|
```

```
|5|nnnnnnnnn|
```

```
|6|qqqqqqqqq|
```

```
|7|ooooooooo|
```

```
|8|yyyyyyyyy|
```

```
|9|yyyyyyyyy|
```

```
|10|vvvvvvvvv|
```

```
+----+-----+
```

```
10 rows in set(0.00 sec)
```

python py_innodb_page_info.py t.ibd

```
[root@nineyou0-43 mytest]#py_innodb_page_info.py-v t.ibd
```

```
page offset 00000000,page typeFile Space Header
```

```
page offset 00000001,page typeInsert Buffer Bitmap
```

```
page offset 00000002,page typeFile Segment inode
```

```
page offset 00000003,page typeB-tree Node,page level0000
```

```
page offset 00000000,page typeFreshly Allocated Page
```

```
page offset 00000000,page typeFreshly Allocated Page
```

```
Total number of page:6:
```

```
Freshly Allocated Page:2
```

```
Insert Buffer Bitmap:1
```

```
File Space Header:1
```

B-tree Node:1

File Segment inode:1

page offset 3hexdumpt.ibd
0x0000c00016K*3=0xc000

```
0000c000 52 1b 24 00 00 00 03 ff ff ff ff ff ff ff|R.$.....|
0000c010 00 00 00 0a 6a e0 ac 93 45 bf 00 00 00 00 00|...j...E.....|
0000c020 00 00 00 00 00 dc 00 1a 0d c0 80 66 00 00 00|.....f....|
0000c030 0d a5 00 02 00 63 00 64 00 00 00 00 00 00|....c.d.....|
0000c040 00 00 00 00 00 00 00 01 ba 00 00 00 dc 00 00|.....|
0000c050 00 02 00 f2 00 00 00 dc 00 00 00 02 00 32 01 00|.....2..|
0000c060 02 00 1c 69 6e 66 69 6d 75 6d 00 05 00 0b 00 00|...infimum.....|
0000c070 73 75 70 72 65 6d 75 6d 0a 00 00 00 10 00 22 00|supremum.....".|
0000c080 00 00 01 00 00 00 51 6d eb 80 00 00 00 2d 01 10|.....Qm.....-..|
0000c090 64 64 64 64 64 64 64 64 0a 00 00 00 18 00|dddddddddd.....|
0000c0a0 22 00 00 00 02 00 00 51 6d ec 80 00 00 00 2d|".....Qm.....-|
0000c0b0 01 10 68 68 68 68 68 68 68 68 68 68 0a 00 00 00|..hhhhhhhhh....|
0000c0c0 20 00 22 00 00 00 03 00 00 51 6d ed 80 00 00 00|. ".....Qm....|
0000c0d0 00 2d 01 10 62 62 62 62 62 62 62 62 62 0a 00|. -..bbbbbbbbbb..|
0000c0e0 04 00 28 00 22 00 00 04 00 00 51 6d ee 80|..(.".....Qm..|
0000c0f0 00 00 00 2d 01 10 69 69 69 69 69 69 69 69 69|...-..iiiiiiiiii|
0000c100 0a 00 00 00 30 00 22 00 00 00 05 00 00 51 6d|....0.".....Qm|
0000c110 ef 80 00 00 00 2d 01 10 6e 6e 6e 6e 6e 6e 6e|.....-..nnnnnnnn|
0000c120 6e 6e 0a 00 00 00 38 00 22 00 00 00 06 00 00 00|nn....8.".....|
0000c130 51 6d f0 80 00 00 00 2d 01 10 71 71 71 71 71 71|Qm.....-..qqqqqq|
0000c140 71 71 71 71 0a 00 00 00 40 00 22 00 00 00 07 00|qqqq...@.".....|
0000c150 00 00 51 6d f1 80 00 00 00 2d 01 10 6f 6f 6f 6f|..Qm.....-..oooo|
0000c160 6f 6f 6f 6f 6f 6f 0a 00 04 00 48 00 22 00 00 00|ooooo....H."...|
```

0000c170 08 00 00 00 51 6d f2 80 00 00 00 2d 01 10 79 79|....Qm.....-..yy|
0000c180 79 79 79 79 79 79 79 79 0a 00 00 00 50 00 22 00|yyyyyyyy....P..|.|
0000c190 00 00 09 00 00 00 51 6d f3 80 00 00 00 2d 01 10|.....Qm.....-..|
0000c1a0 79 79 79 79 79 79 79 79 0a 00 00 00 58 00|yyyyyyyy....X..|.|
0000c1b0 22 00 00 00 0a 00 00 00 51 6d f4 80 00 00 00 2d|".....Qm.....-|
0000c1c0 01 10 76 76 76 76 76 76 76 76 76 76 76 0a 00 00 00|..vvvvvvvvvv....|
0000c1d0 60 00 22 00 00 00 0b 00 00 00 51 6d f5 80 00 00 00|'.".....Qm....|
0000c1e0 00 2d 01 10 6b 6b 6b 6b 6b 6b 6b 6b 6b 6b 0a 00|...kkkkkkkkkk..|
0000c1f0 04 00 68 00 22 00 00 00 0c 00 00 00 51 6d f6 80|..h.".....Qm..|
.....
0000ffc0 00 00 00 00 00 70 0d 1d 0c 95 0c 0d 0b 85 0a fd|....p.....|
0000ffd0 0a 75 09 ed 09 65 08 dd 08 55 07 cd 07 45 06 bd|.u...e...U...E..|
0000ffe0 06 35 05 ad 05 25 04 9d 04 15 03 8d 03 05 02 7d|.5...%.....}|
0000fff0 01 f5 01 6d 00 e5 00 63 95 ae 5d 39 6a e0 ac 93|...m...c...]9j...|

File Header 38

52 1b 24 00Checksum

00 00 00 030

ff ff ff ff0xffffffff

ff ff ff ff0xffffffff

00 00 00 0a 6a e0 ac 93LSN

45 bf0x45bf

00 00 00 00 00 00 00 00

00 00 00 dcSPACE ID

Page HeaderFile TrailerFile Trailer
File HeaderFile Trailer8

95 ae 5d 39 6a e0 ac 93

95 ae 5d 39ChecksumchecksumFile Header
checksum

6a e0 ac 93File HeaderLSN4

56Page HeaderPage Header

Page Header(56 bytes):

PAGE_N_DIR_SLOTS=0x001a

PAGE_HEAP_TOP=0x0dc0

PAGE_N_HEAP=0x8066

PAGE_FREE=0x0000

PAGE_GARBAGE=0x0000

PAGE_LAST_INSERT=0x0da5

PAGE_DIRECTION=0x0002

PAGE_N_DIRECTION=0x0063

PAGE_N_RECS=0x0064

PAGE_MAX_TRX_ID=0x0000000000000000

PAGE_LEVEL=00 00

PAGE_INDEX_ID=0x000000000000001ba

PAGE_BTR_SEG_LEAF=0x000000dc000000200f2

PAGE_BTR_SEG_TOP=0x000000dc00000020032

PAGE_N_DIR_SLOTS=0x001aPage Directory26
20x0000ffc40x0000fff7

0000ffc0 00 00 00 00 00 70 0d 1d 0c 95 0c 0d 0b 85 0a fd|....p.....|
0000ffd0 0a 75 09 ed 09 65 08 dd 08 55 07 cd 07 45 06 bd|.u...e...U...E..|
0000ffe0 06 35 05 ad 05 25 04 9d 04 15 03 8d 03 05 02 7d|.5...%......}|
0000fff0 01 f5 01 6d 00 e5 00 63 95 ae 5d 39 6a e0 ac 93|...m...c...}9j...|

PAGE_HEAP_TOP=0x0dc0
0xc000+0x0dc0=0xcdc0

0000cdb0 00 00 00 2d 01 10 70 70 70 70 70 70 70 70 70|...-..pppppppppp|
0000cdc0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00|.....|
0000cdd0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00|.....|
0000cde0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00|.....|

PAGE_N_HEAP=0x8066Compact0x0802
Redundant2Infinimun
Supremum0x8066-0x8002=0x64100

PAGE_FREE=0x0000
0

PAGE_GARBAGE=0x00000
0

PAGE_LAST_INSERT=0x0da5
0xc0000+0x0da5=0xcda5

0000cda0 00 03 28 f2 cb 00 00 00 64 00 00 00 51 6e 4e 80|..(....d...QnN.|
0000cdb0 00 00 00 2d 01 10 70 70 70 70 70 70 70 70 70|...-..pppppppppp|
0000cdc0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00|.....|

```
05 00 0b 00 00/*recorder header*/
```


Recorder Header
Page HeaderPAGE_PREVPAGE_NEXT
InnoDB

Page Directory0x0000ffc40x0000fff7
Page Directory

0000ffc0 00 00 00 00 70 0d 1d 0c 95 0c 0d 0b 85 0a fd|....p.....|
0000ffd0 0a 75 09 ed 09 65 08 dd 08 55 07 cd 07 45 06 bd|.u...e...U...E..|
0000ffe0 06 35 05 ad 05 25 04 9d 04 15 03 8d 03 05 02 7d|.5...%.....}|
0000fff0 01 f5 01 6d 00 e5 00 63 95 ae 5d 39 6a e0 ac 93|...m...c...}9j...|

Page Directory200 63
0xc06300 700xc070
InfimumSupremumPage Directory
InnoDB
Recorder Headern_ownedInnoDBa
5Page Directory00 e5
0xc0e5

0000c0e0 04 00 28 00 22 00 00 00 04 00 00 00 51 6d ee 80|..(.....Qm..|
0000c0f0 00 00 00 2d 01 10 69 69 69 69 69 69 69 69 69|...-..iiiiiiiiii|
0000c100 0a 00 00 00 30 00 22 00 00 00 05 00 00 00 51 6d|....0.".....Qm|
0000c110 ef 80 00 00 00 2d 01 10 6e 6e 6e 6e 6e 6e 6e|.....nnnnnnnn|
0000c120 6e 6e 0a 00 00 00 38 00 22 00 00 00 06 00 00 00|nn....8.".....|
0000c130 51 6d f0 80 00 00 00 2d 01 10 71 71 71 71 71 71|Qm.....-..qqqqqq|
0000c140 71 71 71 71 0a 00 00 00 40 00 22 00 00 00 07 00|qqqq....@.".....|

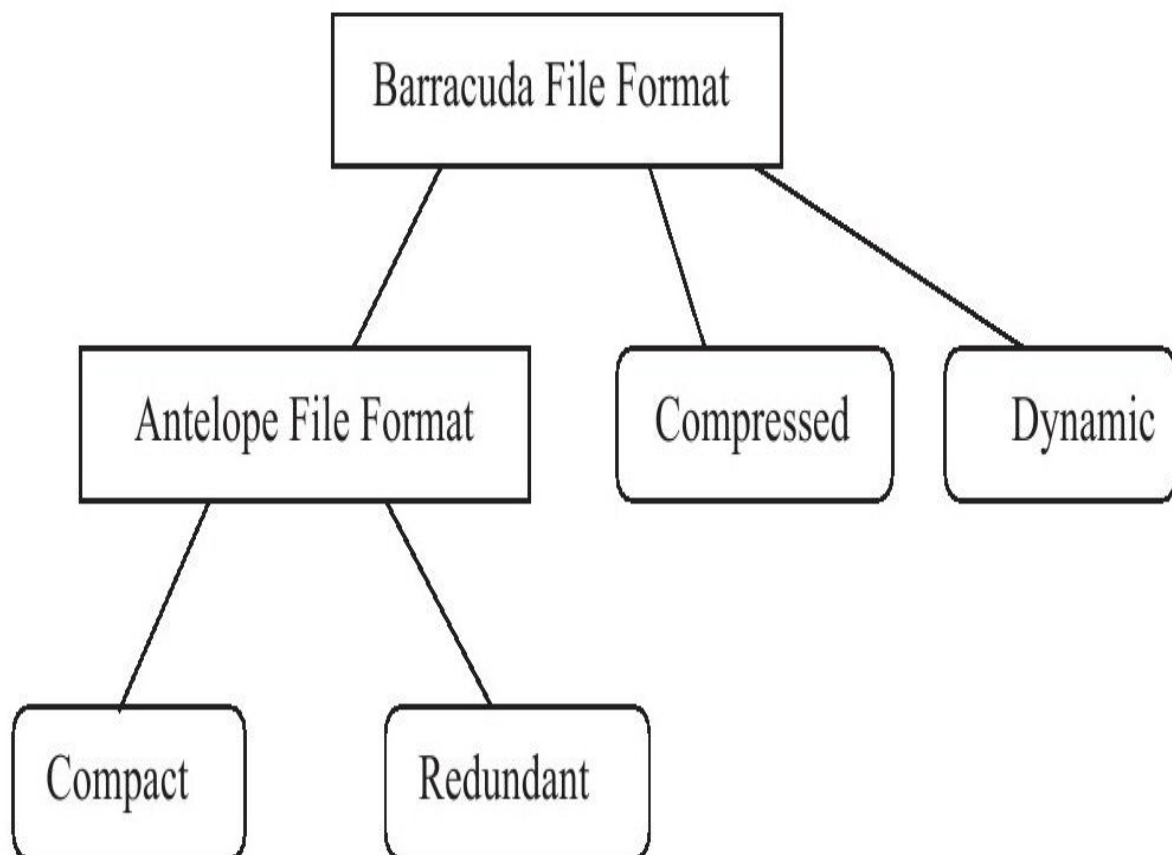
465Record Header04
00 28 00 2248n_owned44
recorder header0x0022
0xc1075

인노디비전 인노DB
인노인사이드

4.5 Named File Formats

InnoDB 1.0.x Off page InnoDB
Named File Formats

InnoDB 1.0.x file format Antelope
Barracuda 4-8
Barracuda Antelope Antelope Compact
Redundant Barracuda Antelope
Compressed Dynamic



InnoDB Plugin Apple InnoDB

/**List of animal names representing file format.*/

static const char*file_format_name_map[]={

"Antelope",

"Barracuda",

"Cheetah",

"Dragon",

"Elk",

"Fox",

"Gazelle",

"Hornet",

"Impala",

"Jaguar",

"Kangaroo",

"Leopard",

"Moose",

"Nautilus",

"Ocelot",

"Porpoise",

"Quail",

"Rabbit",

"Shark",

"Tiger",

"Urchin",

"Viper",

"Whale",

"Xenops",

"Yak",

"Zebra"

};

innodb_file_format innodb InnoDB
innodb

mysql>SELECT @@version\G;

*****1.row*****

@@version:5.1.37

1 row in set(0.00 sec)

mysql>SHOW VARIABLES LIKE 'innodb_version'\G;

*****1.row*****

Variable_name:innodb_version

Value:1.0.4

1 row in set(0.00 sec)

mysql>SHOW VARIABLES LIKE 'innodb_file_format'\G;

*****1.row*****

Variable_name:innodb_file_format

Value:Barracuda

1 row in set(0.00 sec)

innodb_file_format_check innodb InnoDB
ON

InnoDB:Warning:the system tablespace is in a

file format that this version doesn't support

4.6 约束

4.6.1 数据库约束

数据库约束是指对数据库中的数据完整性进行限制和控制的规则。在InnoDB中，约束是通过constraint来定义的。constraint可以限制数据的完整性，确保数据的一致性和准确性。

InnoDB支持多种约束，包括Primary Key（主键）、Unique Key（唯一键）、Foreign Key（外键）、Default（默认值）和NOT NULL（非空）。

InnoDB中，约束是通过constraint来定义的。constraint可以限制数据的完整性，确保数据的一致性和准确性。

约束的定义和修改可以通过SQL语句来完成。

Foreign Key（外键）用于建立两个表之间的关联。

Default（默认值）用于指定字段的默认值。

DEFAULT用于指定字段的默认值。

InnoDB中，约束是通过constraint来定义的。constraint可以限制数据的完整性，确保数据的一致性和准确性。

InnoDB中，约束是通过constraint来定义的。constraint可以限制数据的完整性，确保数据的一致性和准确性。

Primary Key（主键）用于唯一标识表中的每一行。

Unique Key（唯一键）用于确保表中的每一行都是唯一的。

Foreign Key（外键）用于建立两个表之间的关联。

Default（默认值）用于指定字段的默认值。

NOT NULL（非空）用于指定字段不能为空。

PRIMARY
ALTER TABLE

```
mysql> ALTER TABLE u
- ADD UNIQUE KEY uk_id_card(id_card);

Query OK, 0 rows affected (0.19 sec)

Records: 0 Duplicates: 0 Warnings: 0

mysql> SELECT constraint_name, constraint_type
- FROM
- information_schema.TABLE_CONSTRAINTS
- WHERE table_schema='mytest' AND table_name='u';\G;

*****1.row*****

constraint_name: PRIMARY
constraint_type: PRIMARY KEY

*****2.row*****

constraint_name: name
constraint_type: UNIQUE

*****3.row*****

constraint_name: uk_id_card
constraint_type: UNIQUE

3 rows in set (0.00 sec)
```

Foreign Key Foreign Key
p

```
mysql> CREATE TABLE p(
- id INT,
- u_id INT,
- PRIMARY KEY(id),
```

```
- FOREIGN KEY(u_id)REFERENCES p(id));
```

Query OK, 0 rows affected(0.13 sec)

```
mysql>SELECT constraint_name,constraint_type
```

```
- FROM
```

```
- information_schema.TABLE_CONSTRAINTS
```

```
- WHERE table_schema='mytest'and table_name='p'\G;
```

*****1.row*****

constraint_name:PRIMARY

constraint_type:PRIMARY KEY

*****2.row*****

constraint_name:p_ibfk_1

constraint_type:FOREIGN KEY

2 rows in set(0.00 sec)

information_schemaTABLE_CONSTRAINTS
MySQLForeign Key
REFERENTIAL_CONSTRAINTS

```
mysql>SELECT*FROM
```

```
- information_schema.REFERENTIAL_CONSTRAINTS
```

```
- WHERE constraint_schema='mytest'\G;
```

*****1.row*****

CONSTRAINT_CATALOG:NULL

CONSTRAINT_SCHEMA:test2

CONSTRAINT_NAME:p_ibfk_1

UNIQUE_CONSTRAINT_CATALOG:NULL

UNIQUE_CONSTRAINT_SCHEMA:test2

UNIQUE_CONSTRAINT_NAME:PRIMARY

MATCH_OPTION:NONE

UPDATE_RULE:RESTRICT

DELETE_RULE:RESTRICT

TABLE_NAME:p

REFERENCED_TABLE_NAME:p

1 row in set(0.00 sec)

4.6.3 主键与唯一键

主键 (Primary Key) 与唯一键 (Unique Key) 是数据库中的两种键类型。主键用于唯一标识表中的每一行记录，而唯一键用于确保表中的某一列或几列的值是唯一的。

主键和唯一键的主要区别在于：主键不允许有空值 (NULL)，而唯一键允许有空值。此外，一个表只能有一个主键，但可以有一个或多个唯一键。

4.6.4 数据库约束

MySQL数据库约束是指对数据库表中的数据进行限制，以保证数据的完整性和准确性。MySQL数据库支持多种约束，包括主键约束、外键约束、唯一性约束、非空约束（NOT NULL）和空值约束（NULL）等。本文将介绍MySQL数据库中的非空约束（NOT NULL）。

```
mysql>CREATE TABLE a(
  -id INT NOT NULL,
  -date DATE NOT NULL);
Query OK,0 rows affected(0.13 sec)

mysql>INSERT INTO a
  -SELECT NULL,'2009-02-30';
Query OK,1 row affected,2 warnings(0.04 sec)

Records:1 Duplicates:0 Warnings:2

mysql>SHOW WARNINGS\G;

*****1.row*****
Level:Warning
Code:1048
Message:Column'id'cannot be null

*****2.row*****
Level:Warning
Code:1265
Message:Data truncated for column'date'at row 1
2 rows in set(0.00 sec)

mysql>SELECT*FROM a\G;

*****1.row*****

id:0
date:0000-00-00
1 row in set(0.00 sec)
```

NOT NULLNULLdate
'2009-02-30'“”MySQLwarning
MySQL
sql_mode

```
mysql>SET sql_mode='STRICT_TRANS_TABLES';
```

```
Query OK, 0 rows affected (0.00 sec)
```

```
mysql>INSERT INTO a
```

```
->SELECT NULL, '2009-02-30';
```

```
ERROR 1048 (23000): Column 'id' cannot be null
```

```
mysql>INSERT INTO a
```

```
->SELECT 1, '2009-02-30';
```

```
ERROR 1292 (22007): Incorrect date value: '2009-02-30' for column 'date' at row 1
```

sql_modeSTRICT_TRANS_TABLESMySQL
sql_mode
MySQL

4.6.5 ENUM SET

MySQL CHECK ENUM SET
male female
ENUM

```
mysql>CREATE TABLE a(
- id INT,
- sex ENUM('male','female'));
Query OK, 0 rows affected (0.12 sec)

mysql>INSERT INTO a
- SELECT 1,'female';
Query OK, 1 row affected (0.03 sec)

Records: 1 Duplicates: 0 Warnings: 0

mysql>INSERT INTO a
- SELECT 2,'bi';
Query OK, 1 row affected, 1 warning (0.03 sec)

Records: 1 Duplicates: 0 Warnings: 1
```

CHECK
sql_mode

```
mysql>SET sql_mode='STRICT_TRANS_TABLES';
Query OK, 0 rows affected (0.00 sec)

mysql>INSERT INTO a
- SELECT 2,'bi';

ERROR 1265 (01000):Data truncated for column'sex'at row 1
```

enum CHECK
enum SET


```
mysql>INSERT INTO usercash
```

```
->SELECT 1,1000;
```

Query OK,1 row affected(0.03 sec)

Records:1 Duplicates:0 Warnings:0

```
mysql>UPDATE usercash
```

```
->SET cash=cash-(-20)WHERE userid=1;
```

Query OK,1 row affected(0.05 sec)

Rows matched:1 Changed:1 Warnings:0

SQL

```
mysql>CREATE TABLE usercash_err_log(
```

```
->userid INT NOT NULL,
```

```
->old_cash INT UNSIGNED NOT NULL,
```

```
->new_cash INT UNSIGNED NOT NULL,
```

```
->user VARCHAR(30),
```

```
->time DATETIME);
```

Query OK,0 rows affected(0.13 sec)

```
mysql>DELIMITER$$
```

Query OK,0 rows affected(0.00 sec)

```
mysql>CREATE TRIGGER tgr_usercash_update BEFORE UPDATE ON usercash
```

```
->FOR EACH ROW
```

```
->BEGIN
```

```
->IF new.cash-old.cash<0 THEN
```

```
->INSERT INTO usercash_err_log
```

```
->SELECT old.userid,old.cash,new.cash,USER(),NOW();
```

```
->SET new.cash=old.cash;
```

```
->END IF;
```

```
-END;
```

```
-$$
```

```
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> DELIMITER $$
```

```
Query OK, 0 rows affected (0.00 sec)
```

```
CREATE TRIGGER tgr_usercash_update BEFORE UPDATE ON usercash
FOR EACH ROW
BEGIN
    INSERT INTO usercash_err_log (userid, cash, err_msg)
    VALUES (NEW.userid, NEW.cash, 'UPDATE usercash: '
    -- SQL
    --
```

```
mysql> DELETE FROM usercash;
```

```
Query OK, 1 row affected (0.02 sec)
```

```
mysql> INSERT INTO usercash
```

```
VALUES (1, 1000);
```

```
Query OK, 1 row affected (0.03 sec)
```

```
Records: 1 Duplicates: 0 Warnings: 0
```

```
mysql> UPDATE usercash
```

```
SET cash=cash-(-20)
```

```
WHERE userid=1;
```

```
Query OK, 0 rows affected (0.02 sec)
```

```
Rows matched: 1 Changed: 0 Warnings: 0
```

```
mysql> SELECT * FROM usercash\G;
```

```
*****1. row*****
```

```
userid:1
```

```
cash:100
```

```
1 row in set (0.00 sec)
```

```
mysql> SELECT * FROM usercash_err_log\G;
```

```
*****1. row*****
```

```
userid:1
```

old_cash:1000

new_cash:1020

user:root@localhost

time:2009-11-06 11:49:49

Message:Column'id'cannot be null

1 row in set(0.00 sec)



usercash_err_log
SQL

4.6.7 外键

MySQL 支持 MyISAM 和 InnoDB 两种存储引擎。MyISAM 不支持外键，而 InnoDB 支持外键。

```
[CONSTRAINT[symbol]] FOREIGN KEY
[index_name](index_col_name,...)
REFERENCES tbl_name(index_col_name,...)
[ON DELETE reference_option]
[ON UPDATE reference_option]
reference_option:
RESTRICT|CASCADE|SET NULL|NO ACTION
```

创建表使用 `CREATE TABLE`，修改表使用 `ALTER TABLE`。

```
mysql> CREATE TABLE parent(
  -id INT NOT NULL,
  -PRIMARY KEY(id)
  -)ENGINE=INNODB;
Query OK, 0 rows affected (0.13 sec)

mysql> CREATE TABLE child(
  -id INT, parent_id INT,
  -FOREIGN KEY(parent_id) REFERENCES parent(id)
  -)ENGINE=INNODB;
Query OK, 0 rows affected (0.16 sec)
```

外键的 `ON DELETE` 和 `ON UPDATE` 选项用于指定当父表中的记录被删除或更新时，子表中的记录应如何处理。

☐ CASCADE

☐ SET NULL

☐ NO ACTION

☐ RESTRICT

CASCADE DELETE UPDATE
DELETE UPDATE SET NULL DELETE UPDATE
NULL NULL NO ACTION
DELETE UPDATE RESTRICT
DELETE UPDATE ON
DELETE ON UPDATE RESTRICT

Oracle deferred check
SQL MySQL immediate
check MySQL NO ACTION RESTRICT

Oracle InnoDB
Microsoft SQL Server
child
parent_id SHOW CREATE TABLE
InnoDB parent_id

```
mysql> SHOW CREATE TABLE child\G;
```

```
*****1. row*****
```

```
Table: child
```

```
Create Table: CREATE TABLE 'child' (
```

```
  'id' int(11) DEFAULT NULL,
```

```
  'parent_id' int(11) NOT NULL,
```

```
  KEY 'parent_id' ('parent_id'),
```

```
  CONSTRAINT 'child_ibfk_1' FOREIGN KEY ('parent_id') REFERENCES 'parent' ('id')
```

```
) ENGINE=InnoDB DEFAULT CHARSET=utf8
```

1 row in set(0.00 sec)

MySQL

mysql[SET foreign_key_checks=0;

mysql[LOAD DATA.....

mysql[SET foreign_key_checks=1;

4.7 View

MySQL에서 View는 SQL 쿼리를 기반으로 생성된 permanent table과 유사하지만, permanent table과 달리 데이터를 저장하지 않습니다.

4.7.1 View 생성

MySQL에서 View는 base table과 유사하지만, base table과 달리 데이터를 저장하지 않습니다. View는 base table의 데이터를 기반으로 생성되며, base table의 구조와 유사합니다.

MySQL 5.0에서 View를 생성하는 방법은 다음과 같습니다.

```
CREATE  
  
[OR REPLACE]  
  
[ALGORITHM={UNDEFINED|MERGE|TEMPTABLE}]  
  
[DEFINER={user|CURRENT_USER}]  
  
[SQL SECURITY{DEFINER|INVOKER}]  
  
VIEW view_name[(column_list)]  
  
AS select_statement  
  
[WITH[CASCADED|LOCAL]CHECK OPTION]
```

MySQL에서 View를 생성할 때는 updatable view와 non-updatable view로 나뉩니다. updatable view는 WITH CHECK OPTION을 사용하여 생성할 수 있습니다.

```
mysql> CREATE TABLE t(id INT);  
  
Query OK, 0 rows affected (0.13 sec)  
  
mysql> CREATE VIEW v_t  
  
- AS  
  
- SELECT * FROM t WHERE id < 10;
```

Query OK, 0 rows affected (0.00 sec)

```
mysql> INSERT INTO v_t SELECT 20;
```

Query OK, 1 row affected (0.03 sec)

Records: 1 Duplicates: 0 Warnings: 0

```
mysql> SELECT * FROM v_t;
```

Empty set (0.00 sec)

mysql> ALTER VIEW v_t
AS
SELECT * FROM t WHERE id=10
WITH CHECK OPTION;

```
mysql> ALTER VIEW v_t
```

```
- AS
```

```
- SELECT * FROM t WHERE id=10
```

```
- WITH CHECK OPTION;
```

Query OK, 0 rows affected (0.00 sec)

```
mysql> INSERT INTO v_t SELECT 20;
```

ERROR 1369 (HY000): CHECK OPTION failed 'mytest.v_t'

MySQL 5.7.26
MySQL 5.7.26

MySQL 5.7.26 DBA 5.7.26 SHOW TABLES
MySQL 5.7.26 DBA 5.7.26 SHOW TABLES

```
mysql> SHOW TABLES\G;
```

*****1.row*****

Tables_in_mytest:t

*****2.row*****

Tables_in_mytest:v_t

2 rows in set (0.00 sec)

```
mysql> SHOW TABLES FROM test.v_t
mysql> information_schema.TABLES WHERE table_type='BASE TABLE'
mysql>
```

```
mysql> SELECT * FROM information_schema.TABLES
```

```
- WHERE table_type='BASE TABLE'
```

```
- AND table_schema=database()\G;
```

```
*****1.row*****
```

```
TABLE_CATALOG:NULL
```

```
TABLE_SCHEMA:mytest
```

```
TABLE_NAME:t
```

```
TABLE_TYPE:BASE TABLE
```

```
ENGINE:InnoDB
```

```
VERSION:10
```

```
ROW_FORMAT:Compact
```

```
TABLE_ROWS:1
```

```
AVG_ROW_LENGTH:16384
```

```
DATA_LENGTH:16384
```

```
MAX_DATA_LENGTH:0
```

```
INDEX_LENGTH:0
```

```
DATA_FREE:0
```

```
AUTO_INCREMENT:NULL
```

```
CREATE_TIME:2009-11-09 16:27:52
```

```
UPDATE_TIME:NULL
```

```
CHECK_TIME:NULL
```

```
TABLE_COLLATION:utf8_general_ci
```

```
CHECKSUM:NULL
```

```
CREATE_OPTIONS:
```

```
TABLE_COMMENT:
```

1 row in set (0.00 sec)

meta data information_schema VIEWS definer VIEWS

mysql> SELECT * FROM

- information_schema.VIEWS

- WHERE table_schema=database()\G;

*****1. row*****

TABLE_CATALOG: NULL

TABLE_SCHEMA: mytest

TABLE_NAME: v_t

VIEW_DEFINITION: select 'mytest'. 't'. 'id' AS 'id' from 'mytest'. 't' where ('mytest'. 't'. 'id' < 10)

CHECK_OPTION: CASCADED

IS_UPDATABLE: YES

DEFINER: root@localhost

SECURITY_TYPE: DEFINER

CHARACTER_SET_CLIENT: latin1

COLLATION_CONNECTION: latin1_swedish_ci

1 row in set (0.00 sec)

4.7.2 〇〇〇〇

Oracle ——— JOIN GROUP BY SQL Microsoft SQL Server

Oracle

☐ BUILD IMMEDIATE☐ BUILD DEFERRED

BUILD IMMEDIATE [] [] [] [] [] [] [] [] [] [] [] [] [] [] [] [] BUILD
DEFERRED [] [] [] [] [] [] [] [] [] [] [] [] [] [] [] []

SQL

DML

☐ ON DEMAND

ON COMMIT

ON DEMAND ON COMMIT
DDL DML

[illegible]

☐ FAST

☐ COMPLETE

□ FORCE

☐ NEVER

```
FASTXXXXXXXXXXXXXXXXXXXXXXXXXXXXCOMPLETEXXXXXXXXXX
XXXXXXXXXXXXFORCEXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
FASTXXXXXXCOMPLETEXXXNEVERXXXXXXXXXXXXXXXXXXXX
```

```
MySQL[ ]MySQL[ ]
[ ]ON DEMAND[ ]
[ ]
```

```
mysql>CREATE TABLE Orders
- (
-order_id INT UNSIGNED NOT NULL AUTO_INCREMENT,
-product_name VARCHAR(30)NOT NULL,
-price DECIMAL(8,2)NOT NULL,
-amount SMALLINT NOT NULL,
-PRIMARY KEY(order_id)
-ENGINE=InnoDB;
```

Query OK, 0 rows affected (0.13 sec)

```
mysql> INSERT INTO Orders VALUES
```

```
-␣(NULL, 'CPU', 135.5, 1),  
-␣(NULL, 'Memory', 48.2, 3),  
-␣(NULL, 'CPU', 125.6, 3),  
-␣(NULL, 'CPU', 105.3, 4)
```

-□;

Query OK, 4 rows affected (0.03 sec)

Records:4 Duplicates:0 Warnings:0

```
mysql> SELECT * FROM Orders\G;
```

```
*****1.row*****
```

```
order id:1
```

```
product_name:CPU
```

Query OK, 0 rows affected (0.13 sec)

```
mysql>INSERT INTO Orders_MV

-SELECT product_name

-,SUM(price),SUM(amount),AVG(price)

-,COUNT(*)

-FROM Orders

-GROUP BY product_name;
```

Query OK,2 rows affected(0.02 sec)

Records:2 Duplicates:0 Warnings:0

mysql>

mysql>

```
mysql>SELECT*FROM Orders_MV\G;
```

*****1.row*****

product_name:CPU

price:366.40

amount_sum:8

price_avg:122.133

orders_cnt:3

*****2.row*****

product_name:Memory

price:48.20

amount_sum:3

price_avg:48.2

orders_cnt:1

2 rows in set(0.00 sec)

Orders_MV_MVDBA
ON DEMAND
COMPLETEFAST
order_id

Oracle
MySQL
Orders

```
DELIMITER$$

CREATE TRIGGER tgr_Orders_insert

AFTER INSERT ON Orders

FOR EACH ROW

BEGIN

SET@old_price_sum=0;

SET@old_amount_sum=0;

SET@old_price_avg=0;

SET@old_orders_cnt=0;

SELECT IFNULL(price_sum,0),IFNULL(amount_sum,0),IFNULL(price_avg,0),IFNULL(orders_cnt,0)

FROM Orders_MV

WHERE product_name=NEW.product_name

INTO@old_price_sum,@old_amount_sum,@old_price_avg,@old_orders_cnt;

SET@new_price_sum=@old_price_sum+NEW.price;

SET@new_amount_sum=@old_amount_sum+NEW.amount;

SET@new_orders_cnt=@old_orders_cnt+1;

SET@new_price_avg=@new_price_sum/@new_orders_cnt;

REPLACE INTO Orders_MV

VALUES(NEW.product_name,@new_price_sum,@new_amount_sum,@new_price_avg,@new_orders_cnt);

END;

$$

DELIMITER;
```

INSERT
Orders_MV
Orders_MV

```
mysql>INSERT INTO Orders VALUES(NULL,'SSD',299,3);
```

```
Query OK,1 row affected,1 warning(0.03 sec)
```

```
mysql>INSERT INTO Orders VALUES(NULL,'Memory',47.9,5);
```

```
Query OK,1 row affected(0.03 sec)
```

```
mysql>SELECT*FROM Orders_MV\G;
```

```
*****1.row*****
```

```
product_name:CPU
```

```
price:366.40
```

```
amount_sum:8
```

```
price_avg:122.133
```

```
orders_cnt:3
```

```
*****2.row*****
```

```
product_name:Memory
```

```
price:96.10
```

```
amount_sum:8
```

```
price_avg:48.05
```

```
orders_cnt:2
```

```
*****3.row*****
```

```
product_name:SSD
```

```
price:299.00
```

```
amount_sum:3
```

```
price_avg:299
```

```
orders_cnt:1
```

```
3 rows in set(0.00 sec)
```

Orders_MV에 대해 ON_COMMIT UPDATE DELETE DELETE UPDATE
SQL을 실행하면 Orders 테이블에
ON_COMMIT UPDATE DELETE DELETE UPDATE
가 적용됩니다.

MySQLMySQL
Query Rewrite
MySQL

4.8 数据库

4.8.1 数据库引擎

MySQL数据库支持多种存储引擎，包括InnoDB、MyISAM、NDB、CSV、FEDERATED、MERGE等。其中，InnoDB和MyISAM是最常用的两种引擎。

MySQL 5.1版本引入了InnoDB引擎，并使其成为默认引擎。InnoDB引擎支持事务、行级锁、崩溃恢复等功能，适用于需要高可靠性和数据一致性的场景。MyISAM引擎则支持全文索引，适用于读多写少的场景。

MySQL 5.1版本还引入了NDB引擎，这是一种分布式存储引擎，支持高可用性和可扩展性。NDB引擎适用于需要高可用性和可扩展性的场景。此外，MySQL还支持CSV、FEDERATED和MERGE等引擎，以满足不同的需求。

以下命令用于查看MySQL的存储引擎支持情况：

```
mysql> SHOW VARIABLES LIKE '%partition%'\G;

*****1.row*****

Variable_name:have_partitioning

Value:YES

1 row in set(0.00 sec)
```

以下命令用于查看MySQL的存储引擎支持情况：

```
mysql> SHOW PLUGINS\G;

.....

*****2.row*****

Name:partition

Status:ACTIVE

Type:STORAGE ENGINE
```

```
9 rows in set (0.01 sec)
```

[illegible]

```
mysql> CREATE TABLE t1(
-> col1 INT NOT NULL,
-> col2 DATE NOT NULL,
-> col3 INT NOT NULL,
-> col4 INT NOT NULL,
-> UNIQUE KEY(col1,col2)
-> )
->
-> PARTITION BY HASH(col3)
-> PARTITIONS 4;
```

ERROR 1503(HY000):A PRIMARY KEY must include all columns in the table's partitioning function

mysql>CREATE TABLE t1(
col1 INT NULL,
col2 DATE NULL,
col3 INT NULL,
col4 INT NULL,
UNIQUE KEY(col1,col2,col3,col4)
)
PARTITION BY HASH(col3)
PARTITIONS 4;

mysql>CREATE TABLE t1(

col1 INT NULL,

col2 DATE NULL,

col3 INT NULL,

col4 INT NULL,

UNIQUE KEY(col1,col2,col3,col4)

)

PARTITION BY HASH(col3)

PARTITIONS 4;

Query OK,0 rows affected(0.53 sec)

mysql>CREATE TABLE t1(
col1 INT NULL,
col2 DATE NULL,
col3 INT NULL,
col4 INT NULL
)engine=innodb
PARTITION BY HASH(col3)
PARTITIONS 4;
CREATE TABLE t1(
col1 INT NULL,
col2 DATE NULL,

CREATE TABLE t1(

col1 INT NULL,

col2 DATE NULL,

col3 INT NULL,

col4 INT NULL

)engine=innodb

PARTITION BY HASH(col3)

PARTITIONS 4;

CREATE TABLE t1(

col1 INT NULL,

col2 DATE NULL,

```
col3 INT NULL,
```

col4 INT NULL,

```
key(col4)
```

```
)engine=innodb
```

PARTITION BY HASH(col3)

PARTITIONS 4;

[1] _____

[2] _____

4.8.2 RANGE

1. RANGE

MySQL 5.6 版本开始支持 RANGE 分区。创建表 t，id 为 INT 类型，范围 10 到 20 分为 p0，范围 20 到 30 分为 p1。

```
CREATE TABLE t(  
  id INT  
  )ENGINE=INNODB  
  PARTITION BY RANGE(id)(  
    PARTITION p0 VALUES LESS THAN(10),  
    PARTITION p1 VALUES LESS THAN(20));
```

查看表 t 的分区信息，可以看到两个分区 p0 和 p1，分别对应 id 的范围 10 到 20 和 20 到 30。

```
mysql>system ls-lh/usr/local/mysql/data/test2/t*  
-rw-rw-r--1 mysql mysql 8.4K 7月31 14:11/usr/local/mysql/data/test2/t.frm  
-rw-rw-r--1 mysql mysql 28 7月31 14:11/usr/local/mysql/data/test2/t.par  
-rw-rw-r--1 mysql mysql 96K 7月31 14:12/usr/local/mysql/data/test2/t#P#p0.ibd  
-rw-rw-r--1 mysql mysql 96K 7月31 14:12/usr/local/mysql/data/test2/t#P#p1.ibd
```

向表 t 中插入数据。

```
mysql>INSERT INTO t SELECT 9;  
Query OK,1 row affected(0.03 sec)  
Records:1 Duplicates:0 Warnings:0  
mysql>INSERT INTO tSELECT 10;
```


Query OK, 1 row affected (0.03 sec)

Records: 1 Duplicates: 0 Warnings: 0

mysql> INSERT INTO t SELECT 15;

Query OK, 1 row affected (0.03 sec)

Records: 1 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM information_schema.PARTITIONS
WHERE table_schema=database() AND table_name='t';

mysql> SELECT * FROM information_schema.PARTITIONS

WHERE table_schema=database() AND table_name='t';

*****1, row*****

TABLE_CATALOG: NULL

TABLE_SCHEMA: test2

TABLE_NAME: t

PARTITION_NAME: p0

SUBPARTITION_NAME: NULL

PARTITION_ORDINAL_POSITION: 1

SUBPARTITION_ORDINAL_POSITION: NULL

PARTITION_METHOD: RANGE

SUBPARTITION_METHOD: NULL

PARTITION_EXPRESSION: id

SUBPARTITION_EXPRESSION: NULL

PARTITION_DESCRIPTION: 10

TABLE_ROWS: 1

AVG_ROW_LENGTH: 16384

DATA_LENGTH: 16384

MAX_DATA_LENGTH: NULL

INDEX_LENGTH: 0

DATA_FREE: 0

CREATE_TIME:NULL

UPDATE_TIME:NULL

CHECK_TIME:NULL

CHECKSUM:NULL

PARTITION_COMMENT:

NODEGROUP:default

TABLESPACE_NAME:NULL

*****2.row*****

TABLE_CATALOG:NULL

TABLE_SCHEMA:test2

TABLE_NAME:t

PARTITION_NAME:p1

SUBPARTITION_NAME:NULL

PARTITION_ORDINAL_POSITION:2

SUBPARTITION_ORDINAL_POSITION:NULL

PARTITION_METHOD:RANGE

SUBPARTITION_METHOD:NULL

PARTITION_EXPRESSION:id

SUBPARTITION_EXPRESSION:NULL

PARTITION_DESCRIPTION:20

TABLE_ROWS:2

AVG_ROW_LENGTH:8192

DATA_LENGTH:16384

MAX_DATA_LENGTH:NULL

INDEX_LENGTH:0

DATA_FREE:0

CREATE_TIME:NULL

UPDATE_TIME:NULL

CHECK_TIME:NULL

CHECKSUM:NULL

PARTITION_COMMENT:

NODEGROUP:default

TABLESPACE_NAME:NULL

2 rows in set(0.00 sec)

TABLE_ROWS 9 10 15
PARTITION_METHOD RANGE

MySQL 30

mysql>INSERT INTO t SELECT 30;

ERROR 1526(HY000):Table has no partition for value 30

MAXVALUE MAXVALUE
20 MAXVALUE p2

mysql>ALTER TABLE t

->ADD PARTITION(

->partition p2 values less than maxvalue);

Query OK,0 rows affected(0.45 sec)

Records:0 Duplicates:0 Warnings:0

mysql>INSERT INTO t SELECT 30;

Query OK,1 row affected(0.03 sec)

Records:1 Duplicates:0 Warnings:0

RANGE
sales

mysql>CREATE TABLE sales(

```

- money INT UNSIGNED NOT NULL,

- date DATETIME

-)ENGINE=INNODB

- PARTITION by RANGE(YEAR(date)) (

- PARTITION p2008 VALUE LESS THEN(2009),

- PARTITION p2009 VALUE LESS THEN(2010),

- PARTITION p2010 VALUE LESS THEN(2011)

- );

```

Query OK,0 rows affected(0.34 sec)

```
mysql>INSERT INTO sales SELECT 100,'2008-01-01';
```

Query OK,1 row affected(0.03 sec)

Records:1 Duplicates:0 Warnings:0

```
mysql>INSERT INTO sales SELECT 100,'2008-02-01';
```

Query OK,1 row affected(0.03 sec)

Records:1 Duplicates:0 Warnings:0

```
mysql>INSERT INTO sales SELECT 200,'2008-01-02';
```

Query OK,1 row affected(0.04 sec)

Records:1 Duplicates:0 Warnings:0

```
mysql>INSERT INTO sales SELECT 100,'2009-03-01';
```

Query OK,1 row affected(0.03 sec)

Records:1 Duplicates:0 Warnings:0

```
mysql>INSERT INTO sales SELECT 200,'2010-03-01';
```

Query OK,1 row affected(0.03 sec)

Records:1 Duplicates:0 Warnings:0

```

mysql>DELETE
FROM sales WHERE date<='2008-01-01'and date>'2009-01-
01'
mysql>

```

```
mysql>alter table sales drop partition p2008;
```

Query OK, 0 rows affected (0.18 sec)

Records: 0 Duplicates: 0 Warnings: 0

2008

mysql> EXPLAIN PARTITIONS

-> SELECT * FROM sales

-> WHERE date = '2008-01-01' AND date < '2008-12-31' \G;

*****1. row*****

id: 1

select_type: SIMPLE

table: sales

partitions: p2008

type: ALL

possible_keys: NULL

key: NULL

key_len: NULL

ref: NULL

rows: 5

Extra: Using where

1 row in set (0.00 sec)

EXPLAIN PARTITION SQL
p2008 — Partition Pruning

mysql> EXPLAIN PARTITIONS

-> SELECT * FROM sales

-> WHERE date = '2008-01-01' AND date < '2009-01-01' \G;

*****1.row*****

id:1

select_type:SIMPLE

table:sales

partitions:p2008,p2009

type:ALL

possible_keys:NULL

key:NULL

key_len:NULL

ref:NULL

rows:5

Extra:Using where

1 row in set(0.00 sec)

查询日期在'2009-01-01'之前日期='2008-12-31'之前
p2008p2009
SQL

sales

mysqlCREATE TABLE sales(

-money INT UNSIGNED NOT NULL,

-date DATETIME

-)ENGINE=INNODB

-PARTITION by RANGE(YEAR(date)*100+MONTH(date))(

-PARTITION p201001 VALUES LESS THEN(201002),

-PARTITION p201002 VALUES LESS THEN(201003),

-PARTITION p201003 VALUES LESS THEN(201004)

-);

Query OK,0 rows affected(0.37 sec)

MySQL 分区表 SQL 语句

```
mysql>EXPLAIN PARTITIONS
```

```
->SELECT*FROM sales
```

```
->WHERE date='2010-01-01'AND date='2010-01-31'\G;
```

```
*****1.row*****
```

```
id:1
```

```
select_type:SIMPLE
```

```
table:sales
```

```
partitions:p201001,p201002,p201003
```

```
type:ALL
```

```
possible_keys:NULL
```

```
key:NULL
```

```
key_len:NULL
```

```
ref:NULL
```

```
rows:4
```

```
Extra:Using where
```

```
1 row in set(0.00 sec)
```

p201001 p201002 p201003
RANGE YEAR() TO_DAYS()
TO_SECONDS() UNIX_TIMESTAMP()
TO_DAYS()

```
mysql>CREATE TABLE sales(
```

```
->money INT UNSIGNED NOT NULL,
```

```
->date DATETIME
```

```
->)ENGINE=INNODB
```

```
->PARTITION by range(TO_DAYS(date)) (
```

Query OK, 0 rows affected (0.36 sec)

mysql EXPLAIN PARTITIONS

id:1

```
select_type:SIMPLE
```

partitions:p201001

type:ALL

```
possible_keys:NULL
```

key:NULL

```
key_len:NULL
```

ref:NULL

```
rows:4
```

Extra:Using where

```
1 row in set(0.00 sec)
```

2.LIST

LIST RANGE PARTITIONING

```
mysql>CREATE TABLE t(
- a INT,
- b INT)ENGINE=INNODB
- PARTITION BY LIST(b)(
- PARTITION p0 VALUES IN(1,3,5,7,9),
- PARTITION p1 VALUES IN(0,2,4,6,8)
- );
Query OK,0 rows affected(0.26 sec)
```

RANGE PARTITIONING VALUES LESS THAN LIST PARTITIONING VALUES IN

```
mysql>INSERT INTO t SELECT 1,1;
Query OK,1 row affected(0.03 sec)
Records:1 Duplicates:0 Warnings:0
mysql>INSERT INTO t SELECT 1,2;
Query OK,1 row affected(0.03 sec)
Records:1 Duplicates:0 Warnings:0
mysql>INSERT INTO t SELECT 1,3;
Query OK,1 row affected(0.03 sec)
Records:1 Duplicates:0 Warnings:0
mysql>INSERT INTO t SELECT 1,4;
Query OK,1 row affected(0.03 sec)
Records:1 Duplicates:0 Warnings:0
mysql>SELECT table_name,partition_name,table_rows
- FROM information_schema.PARTITIONS
- WHERE table_name='t'AND table_schema=DATABASE()\G;
*****1.row*****
```

```
table_name:t

partition_name:p0

table_rows:2

*****2,row*****

table_name:t

partition_name:p1

table_rows:2

2 rows in set(0.00 sec)
```

MySQL

```
mysql>INSERT INTO t SELECT 1,10;

ERROR 1526(HY000):Table has no partition for value 10
```

INSERTMySQLInnoDB
MySQLInnoDB
MySQL

```
mysql>CREATE TABLE t(

-a INT,

-b INT)ENGINE=MyISAM

-PARTITION BY LIST(b)(

-PARTITION p0 VALUES IN(1,3,5,7,9),

-PARTITION p1 VALUES IN(0,2,4,6,8)

-);

Query OK,0 rows affected(0.05 sec)

mysql>INSERT INTO t VALUES(1,2),(2,4),(6,10),(5,3);

ERROR 1526(HY000):Table has no partition for value 10

mysql>SELECT*FROM t;

+-----+
```

[a|b]

+-----+-----+

[1|2]

[2|4]

+-----+-----+

2 rows in set(0.00 sec)

610531224
InnoDB

mysql>TRUNCATE TABLE t;

Query OK,2 rows affected(0.00 sec)

mysql>ALTER TABLE t ENGINE=InnoDB;

Query OK,0 rows affected(0.25 sec)

Records:0 Duplicates:0 Warnings:0

mysql>INSERT INTO t VALUES(1,2),(2,4),(6,10),(5,3);

ERROR 1526(HY000):Table has no partition for value 10

mysql>SELECT*FROM t;

Empty set(0.00 sec)

610t

3.HASH

HASH
RANGE LIST
MySQL

HASH
CREATE TABLE“PARTITION BY
HASH(expr”“expr”

MySQLのパーティションの作成方法
“PARTITIONS num”
num
PARTITIONS
1

HASH
t
b

```
CREATE TABLE t_hash(  
  
a INT,  
  
b DATETIME  
  
)ENGINE=InnoDB  
  
PARTITION BY HASH(YEAR(b))  
  
PARTITIONS 4;
```

b2010-04-01
t_hash

```
MOD(YEAR('2010-04-01'))4  
  
=MOD(2010,4)  
  
=2
```

p2

```
mysql>INSERT INTO t_hash SELECT 1,'2010-04-01';
```

Query OK,1 row affected(0.04 sec)

Records:1 Duplicates:0 Warnings:0

```
mysql>SELECT table_name,partition_name,table_rows
```

```
-FROM information_schema.PARTITIONS
```

```
-WHERE table_schema=DATABASE()AND table_name='t_hash'\G;
```

```
*****1.row*****
```

table_name:t_hash

partition_name:p0

```

table_rows:0

*****2.row*****

table_name:t_hash

partition_name:p1

table_rows:0

*****3.row*****

table_name:t_hash

partition_name:p2

table_rows:1

*****4.row*****

table_name:t_hash

partition_name:p3

table_rows:0

4 rows in set(0.00 sec)

```

p2 1
 YEAR HASH
 1

MySQL LINEAR HASH
 HASH HASH LINEAR
 HASH LINEAR HASH t_linear_hash t_hash
 1

```

CREATE TABLE t_linear_hash(

a INT,

b DATETIME

)ENGINE=InnoDB

PARTITION BY LINEAR HASH(YEAR(b))

PARTITIONS 4;

```

```
partition_name:p3
```

table_rows:0

4 rows in set(0.01 sec)

LINEAR HASH
HASH

4.KEY

KEY
MySQL
NDB Cluster
MySQL
MD5
PASSWORD()

```
mysql>CREATE TABLE t_key(
```

```
-a INT,
```

```
-b DATETIME)ENGINE=InnoDB
```

```
-PARTITION BY KEY(b)
```

```
-PARTITIONS 4;
```

Query OK,0 rows affected(0.43 sec)

KEY
LINEAR
HASH
powers-of-two

5.COLUMNS

RANGE
LIST
HASH
KEY
interger
YEAR()
TO_DAYS()
MONTH()
MySQL5.5
COLUMNS
RANGE
LIST
COLUMNS
RANGECOLUMNS

COLUMNS

INT SMALLINT TINYINT BIGINT FLOAT
DECIMAL

DATE DATETIME

CHAR VARCHAR BINARY VARBINARY BLOB TEXT

YEAR() TO_DAYS()
COLUMNS

```
CREATE TABLE t_columns_range(  
  
  a INT,  
  
  b DATETIME  
  
)ENGINE=INNODB  
  
PARTITION BY RANGE COLUMNS(b)(  
  
PARTITION p0 VALUES LESS THAN('2009-01-01'),  
  
PARTITION p1 VALUES LESS THAN('2010-01-01')  
  
);
```

```
CREATE TABLE customers_1(  
  
  first_name VARCHAR(25),  
  
  last_name VARCHAR(25),  
  
  street_1 VARCHAR(30),  
  
  street_2 VARCHAR(30),  
  
  city VARCHAR(15),  
  
  renewal DATE  
  
)  
  
PARTITION BY LIST COLUMNS(city)(
```


4.8.3 复合分区

复合分区(subpartitioning)是MySQL 5.6引入的一个新特性。它允许在一个分区表中，对每个分区再进行子分区。复合分区可以分为两种：RANGE复合分区和LIST复合分区。此外，还可以结合HASH和KEY函数进行子分区。

```
mysql> CREATE TABLE ts(a INT,b DATE)engine=innodb
PARTITION BY RANGE(YEAR(b))
SUBPARTITION BY HASH(TO_DAYS(b))
SUBPARTITIONS 2(
PARTITION p0 VALUES LESS THAN(1990),
PARTITION p1 VALUES LESS THAN(2000),
PARTITION p2 VALUES LESS THAN MAXVALUE
);
Query OK, 0 rows affected (0.01 sec)

mysql> system ls -lh /usr/local/mysql/data/test2/ts*

-rw-rw---- 1 mysql mysql 8.4K Aug 1 15:50 /usr/local/mysql/data/test2/ts.frm
-rw-rw---- 1 mysql mysql 96 Aug 1 15:50 /usr/local/mysql/data/test2/ts.par
-rw-rw---- 1 mysql mysql 96K Aug 1 15:50 /usr/local/mysql/data/test2/ts#P#p0#SP#p0sp0.ibd
-rw-rw---- 1 mysql mysql 96K Aug 1 15:50 /usr/local/mysql/data/test2/ts#P#p0#SP#p0sp1.ibd
-rw-rw---- 1 mysql mysql 96K Aug 1 15:50 /usr/local/mysql/data/test2/ts#P#p1#SP#p1sp0.ibd
-rw-rw---- 1 mysql mysql 96K Aug 1 15:50 /usr/local/mysql/data/test2/
ts#P#p1#SP#p1sp1.ibd
-rw-rw---- 1 mysql mysql 96K Aug 1 15:50 /usr/local/mysql/data/test2/ts#P#p2#SP#p2sp0.ibd
-rw-rw---- 1 mysql mysql 96K Aug 1 15:50 /usr/local/mysql/data/test2/ts#P#p2#SP#p2sp1.ibd
```

在上面的例子中，我们创建了一个名为ts的表，它按照年份(YEAR(b))进行分区。每个分区(p0, p1, p2)又被进一步划分为两个子分区(sp0, sp1)，这是通过HASH函数实现的。因此，总共有3个分区，每个分区有2个子分区，总共是3×2=6个子分区。在文件系统中，我们可以看到每个子分区对应的ibd文件。

```
mysql>CREATE TABLE ts(a INT,b DATE)
```

```
-PARTITION BY RANGE(YEAR(b))
```

```
-SUBPARTITION BY HASH(TO_DAYS(b))
```

```
-PARTITION p0 VALUES LESS THAN(1990)
```

```
-SUBPARTITION s0,
```

```
-SUBPARTITION s1
```

```
-),
```

```
-PARTITION p1 VALUES LESS THAN(2000)
```

```
-SUBPARTITION s2,
```

```
-SUBPARTITION s3
```

```
-),
```

```
-PARTITION p2 VALUES LESS THAN MAXVALUE
```

```
-SUBPARTITION s4,
```

```
-SUBPARTITION s5
```

```
-)
```

```
;
```

```
Query OK,0 rows affected(0.00 sec)
```

mysql>show full columns from ts;

mysql>show full partitions from ts;

mysql>show full partitions from ts\G

mysql>show full partitions from ts\G

```
mysql>CREATE TABLE ts(a INT,b DATE)
```

```
-PARTITION BY RANGE(YEAR(b))
```

```
-SUBPARTITION BY HASH(TO_DAYS(b))
```

```
-PARTITION p0 VALUES LESS THAN(1990)
```

```
-SUBPARTITION s0,
```

```
-SUBPARTITION s1
```

```

-[]),

-[]PARTITION p1 VALUES LESS THAN(2000),

-[]PARTITION p2 VALUES LESS THAN MAXVALUE(

-[]SUBPARTITION s2,

-[]SUBPARTITION s3

-[])

-[]);

ERROR 1064(42000):Wrong number of subpartitions defined,mismatch with previous setting near'

PARTITION p2 VALUES LESS THAN MAXVALUE(

SUBPARTITION s2,

SUBPARTITION s3

)

)'at line 8

```

[][]SUBPARTITION[][][][][][][][][][]

[][][][][][][][][][][][][][][][][][]

```

mysql[]CREATE TABLE ts(a INT,b DATE)

-[]PARTITION BY RANGE(YEAR(b))

-[]SUBPARTITION BY HASH(TO_DAYS(b)) (

-[]PARTITION p0 VALUES LESS THAN(1990) (

-[]SUBPARTITION s0,

-[]SUBPARTITION s1

-[]),

-[]PARTITION p1 VALUES LESS THAN(2000) (

-[]SUBPARTITION s0,

-[]SUBPARTITION s1

-[]),

-[]PARTITION p2 VALUES LESS THAN MAXVALUE(

```

```
-SUBPARTITION s0,
```

```
-SUBPARTITION s1
```

```
-)
```

```
);
```

```
ERROR 1517(HY000):Duplicate partition name s0
```

6

/disk0/disk1/disk2

```
mysql>CREATE TABLE ts(a INT,b DATE)ENGINE=MYISAM
```

```
-PARTITION BY RANGE(YEAR(b))
```

```
-SUBPARTITION BY HASH(TO_DAYS(b))
```

```
-PARTITION p0 VALUES LESS THAN(2000)
```

```
-SUBPARTITION s0
```

```
-DATA DIRECTORY='/disk0/data'
```

```
-INDEX DIRECTORY='/disk0/idx',
```

```
-SUBPARTITION s1
```

```
-DATA DIRECTORY='/disk1/data'
```

```
-INDEX DIRECTORY='/disk1/idx'
```

```
),
```

```
-PARTITION p1 VALUES LESS THAN(2010)
```

```
-SUBPARTITION s2
```

```
-DATA DIRECTORY='/disk2/data'
```

```
-INDEX DIRECTORY='/disk2/idx',
```

```
-SUBPARTITION s3
```

```
-DATA DIRECTORY='/disk3/data'
```

```
-INDEX DIRECTORY='/disk3/idx'
```

```
),
```

```
-PARTITION p2 VALUES LESS THAN MAXVALUE
```

```
-SUBPARTITION s4
```

```
- DATA DIRECTORY='/disk4/data'

- INDEX DIRECTORY='/disk4/idx',

- SUBPARTITION s5

- DATA DIRECTORY='/disk5/data'

- INDEX DIRECTORY='/disk5/idx'

- )

- );

Query OK, 0 rows affected (0.02 sec)
```

MySQL InnoDB 分区表 分区索引 DATA DIRECTORY INDEX DIRECTORY 目录设置

```
mysql> CREATE TABLE ts(a INT,b DATE)engine=innodb
```

```
- PARTITION BY RANGE(YEAR(b))

- SUBPARTITION BY HASH(TO_DAYS(b)) (

- PARTITION p0 VALUES LESS THAN(2000) (

- SUBPARTITION s0

- DATA DIRECTORY='/disk0/data'

- INDEX DIRECTORY='/disk0/idx',

- SUBPARTITION s1

- DATA DIRECTORY='/disk1/data'

- INDEX DIRECTORY='/disk1/idx'

- ),

- PARTITION p1 VALUES LESS THAN(2010) (

- SUBPARTITION s2

- DATA DIRECTORY='/disk2/data'

- INDEX DIRECTORY='/disk2/idx',

- SUBPARTITION s3

- DATA DIRECTORY='/disk3/data'

- INDEX DIRECTORY='/disk3/idx'
```

```
-[]),  
  
-[]PARTITION p2 VALUES LESS THAN MAXVALUE(  
  
-[]SUBPARTITION s4  
  
-[]DATA DIRECTORY='/disk4/data'  
  
-[]INDEX DIRECTORY='/disk4/idx',  
  
-[]SUBPARTITION s5  
  
-[]DATA DIRECTORY='/disk5/data'  
  
-[]INDEX DIRECTORY='/disk5/idx'  
  
-[])  
  
-[]);
```

Query OK,0 rows affected(0.02 sec)

```
mysql[]system ls-lh/usr/local/mysql/data/test2/ts*
```

```
-rw-rw----1 mysql mysql 8.4K Aug 1 16:24/usr/local/mysql/data/test2/ts.frm  
  
-rw-rw----1 mysql mysql 80 Aug 1 16:24/usr/local/mysql/data/test2/ts.par  
  
-rw-rw----1 mysql mysql 96K Aug 1 16:25/usr/local/mysql/data/test2/ts#P#p0#SP#s0.ibd  
  
-rw-rw----1 mysql mysql 96K Aug 1 16:25/usr/local/mysql/data/test2/ts#P#p0#SP#s1.ibd  
  
-rw-rw----1 mysql mysql 96K Aug 1 16:25/usr/local/mysql/data/test2/ts#P#p1#SP#s2.ibd  
  
-rw-rw----1 mysql mysql 96K Aug 1 16:25/usr/local/mysql/data/test2/ts#P#p1#SP#s3.ibd  
  
-rw-rw----1 mysql mysql 96K Aug 1 16:25/usr/local/mysql/data/test2/ts#P#p2#SP#s4.ibd  
  
-rw-rw----1 mysql mysql 96K Aug 1 16:25/usr/local/mysql/data/test2/ts#P#p2#SP#s5.ibd
```

4.8.4 如何NULL

MySQL如何NULL如何如何如何如何如何如何如何如何MySQL如何如何如何NULL如何如何如何NULL如何MySQL如何如何NULL如何ORDER BY如何如何如何如何如何如何如何MySQL如何如何NULL如何如何如何

如何RANGE如何如何如何如何如何NULL如何MySQL如何如何如何如何如何如何如何如何如何

```
mysql>CREATE TABLE t_range(  
-a INT,  
-b INT)ENGINE=InnoDB  
-PARTITION BY RANGE(b)(  
-PARTITION p0 VALUES LESS THAN(10),  
-PARTITION p1 VALUES LESS THAN(20),  
-PARTITION p2 VALUES LESS THAN MAXVALUE  
-);  
Query OK,0 rows affected(0.01 sec)
```

如何如何如何1,1如何1,NULL如何如何如何如何如何如何如何如何如何

```
mysql>INSERT INTO t_range SELECT 1,1;  
Query OK,1 row affected(0.00 sec)  
Records:1 Duplicates:0 Warnings:0  
mysql>INSERT INTO t_range SELECT 1,NULL;  
Query OK,1 row affected(0.00 sec)  
Records:1 Duplicates:0 Warnings:0  
mysql>SELECT*FROM t_range\G;  
*****1.row*****
```


a:1

b:1

*****2.row*****

a:1

b:NULL

2 rows in set(0.00 sec)

mysql>SELECT table_name,partition_name,table_rows

->FROM information_schema.PARTITIONS

->WHERE table_schema=DATABASE()AND table_name='t_range'\G;

*****1.row*****

table_name:t_range

partition_name:p0

table_rows:2

*****2.row*****

table_name:t_range

partition_name:p1

table_rows:0

*****3.row*****

table_name:t_range

partition_name:p2

table_rows:0

3 rows in set(0.00 sec)

| | | | | |
|----------------|----|------------|-------|------|
| partition_name | p0 | table_rows | RANGE | NULL |
| partition_name | p0 | 10 | NULL | |

mysql>ALTER TABLE t_range DROP PARTITION p0;

Query OK,0 rows affected(0.01 sec)

Records:0 Duplicates:0 Warnings:0

```
mysql>SELECT*FROM t_range;
```

```
Empty set(0.00 sec)
```

LISTNULLNULL

```
mysql>CREATE TABLE t_list(
```

```
-a INT,
```

```
-b INT)ENGINE=INNODB
```

```
-PARTITION BY LIST(b)(
```

```
-PARTITION p0 VALUES IN(1,3,5,7,9),
```

```
-PARTITION p1 VALUES IN(0,2,4,6,8)
```

```
);
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>INSERT INTO t_list SELECT 1,NULL;
```

```
ERROR 1526(HY000):Table has no partition for value NULL
```

p0NULL

```
mysql>CREATE TABLE t_list(
```

```
-a INT,
```

```
-b INT)ENGINE=INNODB
```

```
-PARTITION BY LIST(b)(
```

```
-PARTITION p0 VALUES IN(1,3,5,7,9,NULL),
```

```
-PARTITION p1 VALUES IN(0,2,4,6,8)
```

```
);
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>INSERT INTO t_list SELECT 1,NULL;
```

```
Query OK,1 row affected(0.00 sec)
```

```
Records:1 Duplicates:0 Warnings:0
```

```
mysql>SELECT table_name,partition_name,table_rows
```

```

- FROM information_schema.PARTITIONS

- WHERE table_schema=DATABASE()AND table_name='t_list'\G;

*****1.row*****

table_name:t_list

partition_name:p0

table_rows:1

*****2.row*****

table_name:t_list

partition_name:p1

table_rows:0

2 rows in set(0.00 sec)

```

HASH KEY NULL RANGE LIST

NULL 0

```

mysql>CREATE TABLE t_hash(

-a INT,

-b INT)ENGINE=InnoDB

-PARTITION BY HASH(b)

-PARTITIONS 4;

Query OK,0 rows affected(0.00 sec)

mysql>INSERT INTO t_hash SELECT 1,0;

Query OK,1 row affected(0.00 sec)

Records:1 Duplicates:0 Warnings:0

mysql>INSERT INTO t_hash SELECT 1,NULL;

Query OK,1 row affected(0.01 sec)

Records:1 Duplicates:0 Warnings:0

mysql>SELECT table_name,partition_name,table_rows

- FROM information_schema.PARTITIONS

- WHERE table_schema=DATABASE()AND table_name='t_hash'\G;

```

*****1.row*****

table_name:t_hash

partition_name:p0

table_rows:2

*****2.row*****

table_name:t_hash

partition_name:p1

table_rows:0

*****3.row*****

table_name:t_hash

partition_name:p2

table_rows:0

*****4.row*****

table_name:t_hash

partition_name:p3

table_rows:0

4 rows in set(0.00 sec)

4.8.5

[illegible]

OLTP Blog OLAP OLTP OLAP OLAP

OLAP
1
Partition Pruning

OLTP 10%
B+ B+ 2 3
IO B+

```

1000W
10 HASH
100W
SELECT*FROM TABLE WHERE PK=@pk
100W 1000W B+ 2
1000W B+ 3 100W B+ 2
1 IO
SQL SELECT*FROM TABLE WHERE
KEY=@key KEY 10 2 IO
20 IO KEY 2 3 IO

```

```

ProfileIDHASHHASH10Profile
1000W

```

```
mysql>CREATE TABLE'Profile'(  
- 'id'int(11)NOT NULL AUTO_INCREMENT,  
- 'nickname'varchar(20)NOT NULL DEFAULT'',
```

```

- 'password'varchar(32)NOT NULL DEFAULT'',

- 'sex'char(1)NOT NULL DEFAULT'',

- 'rdate'date NOT NULL DEFAULT'0000-00-00',

- PRIMARY KEY('id'),

- KEY'nickname'('nickname')

- )ENGINE=InnoDB

- PARTITION BY HASH(id)

- PARTITIONS 10;

```

Query OK,0 rows affected(1.29 sec)

```
mysql>SELECT COUNT(nickname)FROM Profile;
```

```
*****1.row*****
```

count(1):9999248

1 row in set(1 min 24.62 sec)

HASH

```
mysql>SELECT table_name,partition_name,table_rows
```

```
- FROM information_schema.PARTITIONS
```

```
- WHERE table_schema=DATABASE()AND table_name='Profile'\G;
```

```
*****1.row*****
```

table_name:Profile

partition_name:p0

table_rows:990703

```
*****2.row*****
```

table_name:Profile

partition_name:p1

table_rows:1086519

```
*****3.row*****
```

table_name:Profile

partition_name:p2

table_rows:976474

*****4.row*****

table_name:Profile

partition_name:p3

table_rows:986937

*****5.row*****

table_name:Profile

partition_name:p4

table_rows:993667

*****6.row*****

table_name:Profile

partition_name:p5

table_rows:978046

*****7.row*****

table_name:Profile

partition_name:p6

table_rows:990703

*****8.row*****

table_name:Profile

partition_name:p7

table_rows:978639

*****9.row*****

table_name:Profile

partition_name:p8

table_rows:1085334

*****10.row*****

table_name:Profile

partition_name:p9

table_rows:982788

10 rows in set(0.80 sec)

mysql EXPLAIN PARTITIONS SELECT * FROM Profile WHERE id=1\G

1 row in set (0.00 sec)

mysql EXPLAIN PARTITIONS SELECT * FROM Profile WHERE id=1\G;

*****1.row*****

id:1

select_type:SIMPLE

table:Profile

partitions:p1

type:const

possible_keys:PRIMARY

key:PRIMARY

key_len:4

ref:const

rows:1

Extra:

1 row in set (0.00 sec)

mysql EXPLAIN PARTITIONS SELECT * FROM Profile WHERE nickname='david'\G

mysql EXPLAIN PARTITIONS

- SELECT * FROM Profile WHERE nickname='david'\G;

*****1.row*****

id:1

select_type:SIMPLE

table:Profile

partitions:p0,p1,p2,p3,p4,p5,p6,p7,p8,p9

type:ref

possible_keys:nickname

key:nickname

key_len:62

ref:const

rows:10

Extra:Using where

1 row in set(0.00 sec)

MySQL

mysql[SELECT*FROM Profile WHERE nickname='david'\G;

*****1.row*****

id:5566

nickname:david

password:3e35d1025659d07ae28e0069ec51ab92

sex:M

rdate:2003-09-20

1 row in set(1.05 sec)

1.05 IO
20 30 SQL 0.26

InnoDB OLTP
OLTP

4.8.6 表空间管理

MySQL 5.6 支持 `ALTER TABLE...EXCHANGE PARTITION` 语句，用于交换表空间。该语句允许用户在不影响表数据的情况下，将表的一个或多个分区移动到另一个表空间。这通常用于维护或迁移数据。

`ALTER TABLE...EXCHANGE PARTITION` 语句的语法如下：

`ALTER TABLE table_name EXCHANGE PARTITION partition_name`

其中，`table_name` 是表的名称，`partition_name` 是要交换的分区名称。

该语句要求被交换的分区必须是主键分区或唯一分区。

此外，还可以使用 `ALTER`、`INSERT`、`CREATE` 和 `DROP` 语句来管理表空间。

例如，可以使用 `ALTER` 语句来修改表空间：

`ALTER TABLE table_name EXCHANGE PARTITION partition_name`

使用 `AUTO_INCREMENT` 属性来指定自动递增的列。

还可以使用 `RANGE` 属性来指定范围，例如：

```
CREATE TABLE e(  
  id INT NOT NULL,  
  fname VARCHAR(30),  
  lname VARCHAR(30)  
)  
  
PARTITION BY RANGE(id)(  
  PARTITION p0 VALUES LESS THAN(50),  
  PARTITION p1 VALUES LESS THAN(100),  
  PARTITION p2 VALUES LESS THAN(150),  
  PARTITION p3 VALUES LESS THAN(MAXVALUE))
```

```
);

INSERT INTO e VALUES

(1669,"Jim","Smith"),

(337,"Mary","Jones"),

(16,"Frank","White"),

(2005,"Linda","Black");
```

mysql>CREATE TABLE e2 LIKE e;

Query OK,0 rows affected(1.34 sec)

mysql>ALTER TABLE e2 REMOVE PARTITIONING;

Query OK,0 rows affected(0.90 sec)

Records:0 Duplicates:0 Warnings:0

mysql>SELECT PARTITION_NAME,TABLE_ROWS

-FROM INFORMATION_SCHEMA.PARTITIONS

-WHERE TABLE_NAME='e';

+-----+-----+

|PARTITION_NAME|TABLE_ROWS|

+-----+-----+

|p0|1|

|p1|0|

|p2|0|

|p3|3|

+-----+-----+

4 rows in set(0.00 sec)

mysql> ALTER TABLE e EXCHANGE PARTITION p0 WITH TABLE e2;

Query OK, 0 rows affected (0.28 sec)

mysql> SELECT PARTITION_NAME, TABLE_ROWS

FROM INFORMATION_SCHEMA.PARTITIONS

WHERE TABLE_NAME='e';

+

|PARTITION_NAME|TABLE_ROWS|

+

|p0|0|

|p1|0|

|p2|0|

|p3|3|

+

4 rows in set (0.00 sec)

mysql> SELECT * FROM e2;

+

|id|fname|lname|

+

|16|Frank|White|

+

1 row in set (0.00 sec)

4.9 〇〇

```

InnoDB
InnoDB
MySQL
InnoDB
MySQL

```

MySQL RANGE LIST HASH KEY COLUMNS
HASH KEY OLTP

5 数据库

数据库是存储和管理数据的系统。数据库系统由数据库、数据库管理系统、数据库应用程序和数据库用户组成。

数据库系统——数据库系统是由数据库、数据库管理系统、数据库应用程序和数据库用户组成的。数据库系统是数据库系统的核心，它负责存储和管理数据。数据库管理系统是数据库系统的核心，它负责管理数据库。数据库应用程序是数据库系统的核心，它负责使用数据库。数据库用户是数据库系统的核心，他们使用数据库系统来存储和管理数据。SQL是数据库系统的核心，它负责管理数据库。DBA是数据库系统的核心，他们负责使用数据库。MySQL是数据库系统的核心，它负责管理数据库。iostat是数据库系统的核心，它负责管理数据库。100%是数据库系统的核心，它负责管理数据库。20%是数据库系统的核心，它负责管理数据库。

InnoDB是MySQL的默认存储引擎。MySQL是数据库系统的核心，它负责管理数据库。MySQL是数据库系统的核心，它负责管理数据库。MySQL是数据库系统的核心，它负责管理数据库。

5.1 InnoDB

InnoDB

$$\square B + \square\square\square$$

□ □ □ □ □

□ □ □ □

□□□□□□InnoDB□□□□□□□□□□□□□□□□InnoDB□□□□□□□□□□
□□□□□□□□□□□□□□□□□□□□□□□□□□□□

[illegible]

B+ binary balance B+
 B+

DBA B+ B+

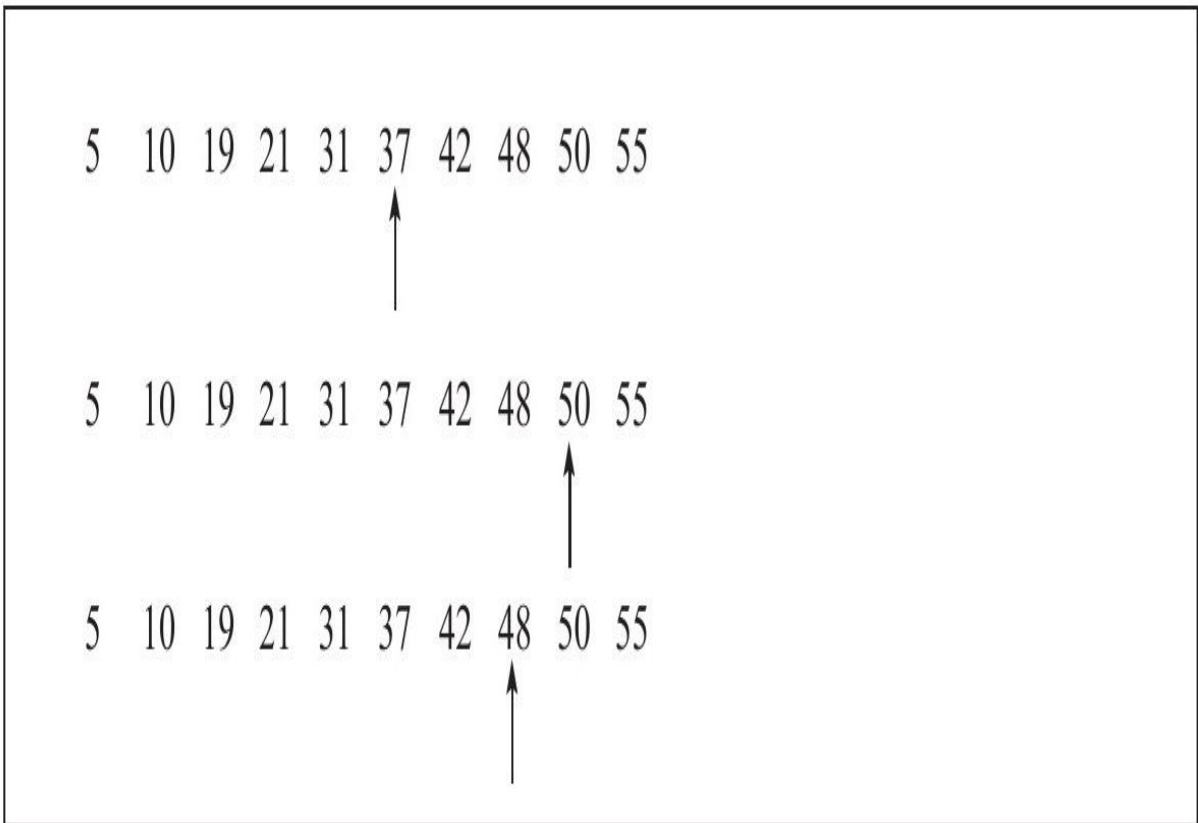
5.2 〇〇〇〇〇〇〇〇

B+
B+

5.2.1 五五五五五

```
#####binary search#####  
#####  
#####  
#####  
#####
```

5 10 19 21 31 37 42 48 50 52 10 10 10 10 10 48
5-1



5-1 3 48 8
 5 1 4
 10 1+2+3+4+5+6+7+8+9+10 / 10 = 5.5
 4+3+2+4+3+1+4+3+2+3 / 10 = 2.9
 10 4

1946
1962
Page Directory
Page Directory

5.2.2 平衡二叉树

平衡二叉树(B+)是一种特殊的二叉树，其特点是任意节点的左右子树的高度差不超过1。这种结构可以保证查找、插入和删除操作的时间复杂度为 $O(\log n)$ 。

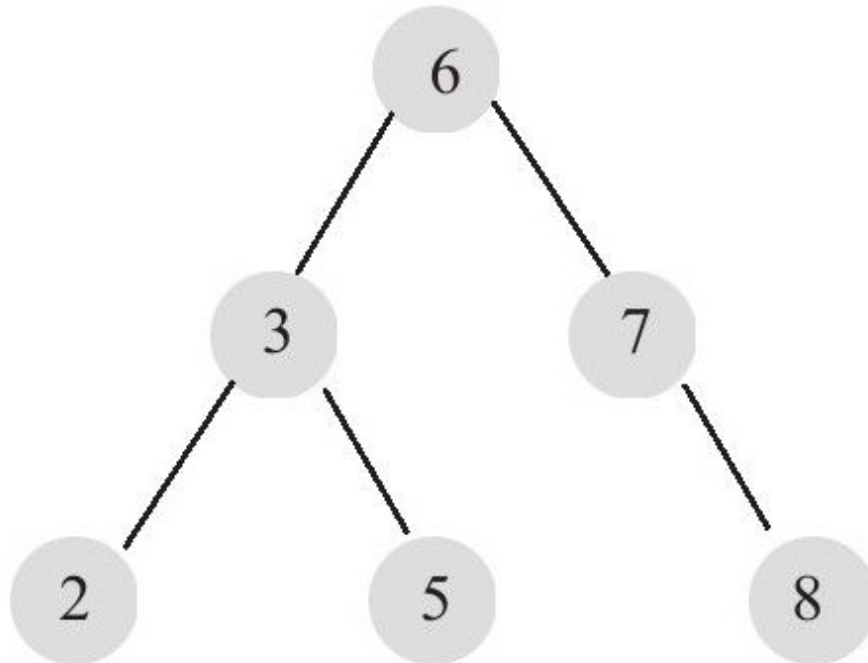


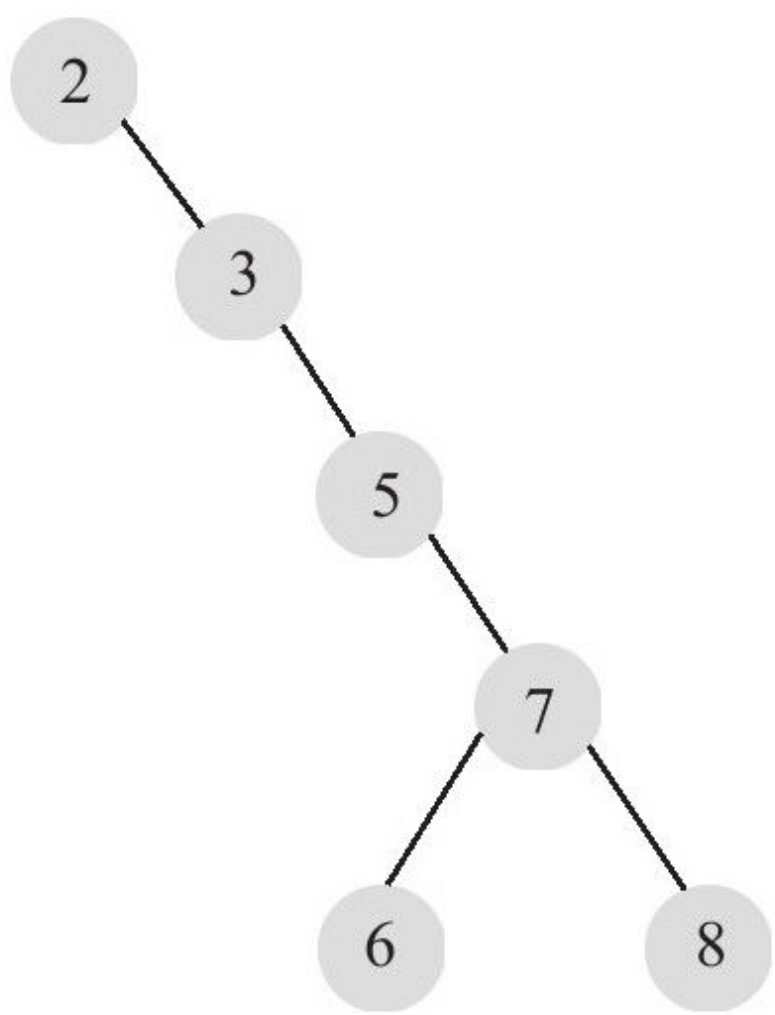
图 5-2 平衡二叉树

图5-2所示的平衡二叉树，其根节点为6，左子树为3，右子树为7。节点3的左子树为2，右子树为5。节点7的右子树为8。该树满足平衡二叉树的定义，即任意节点的左右子树的高度差不超过1。

图5-2所示的平衡二叉树，其根节点为6，左子树为3，右子树为7。节点3的左子树为2，右子树为5。节点7的右子树为8。该树满足平衡二叉树的定义，即任意节点的左右子树的高度差不超过1。该树的平衡因子为1，即左右子树的高度差为1。

该树的平衡因子为1，即左右子树的高度差为1。该树的平衡因子为1，即左右子树的高度差为1。

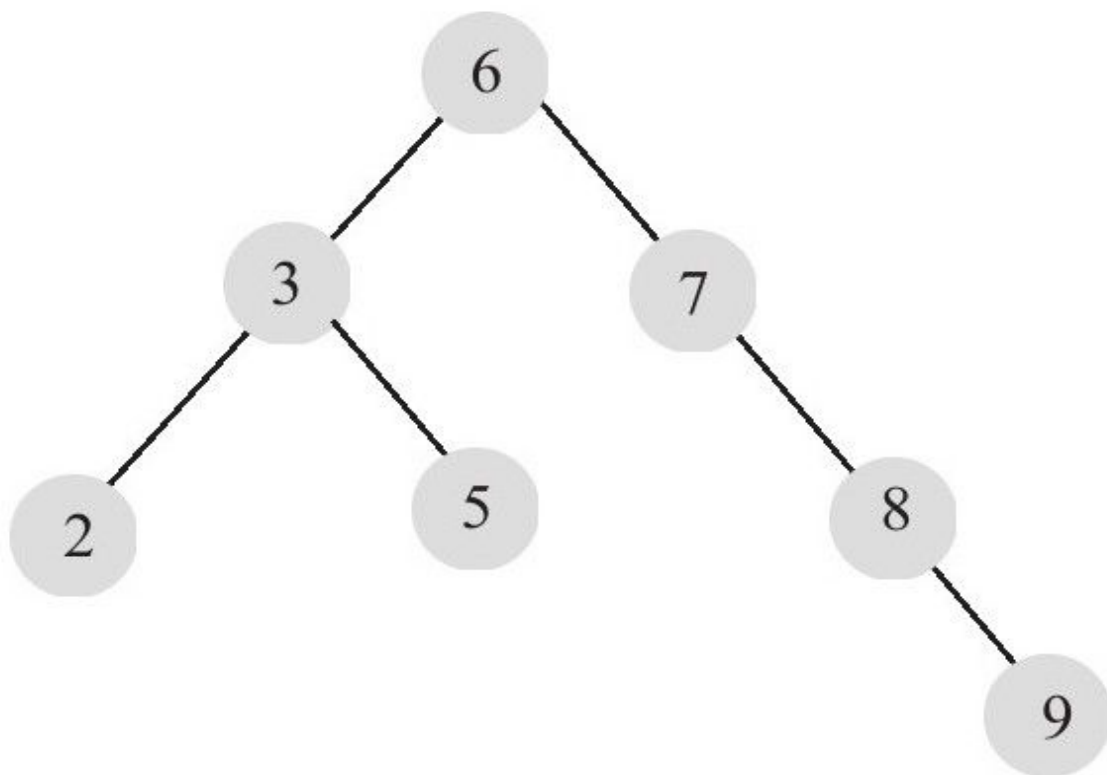
5-3 1+2+3+4+5+5/6=3.16
AVL



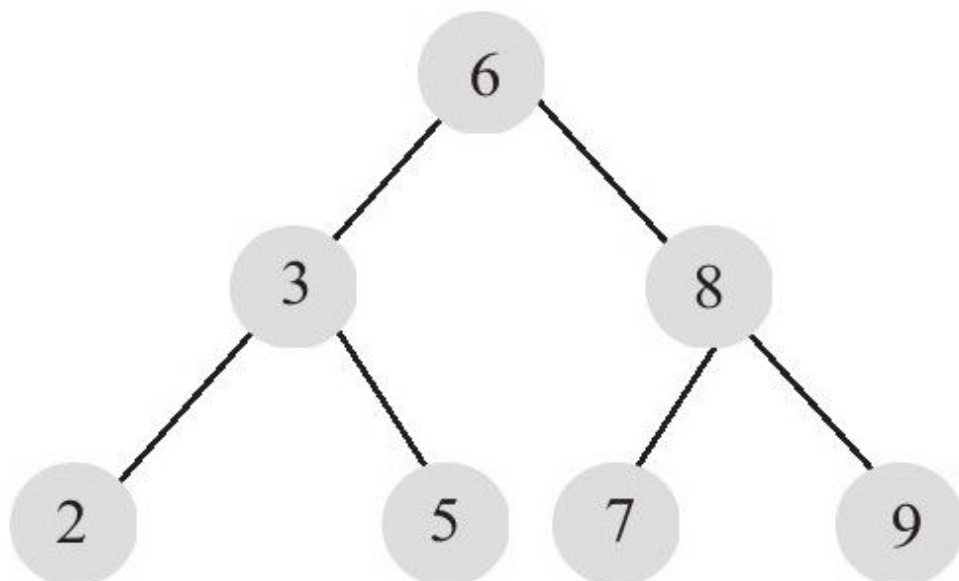
5-3

5-3 5-2 1 9

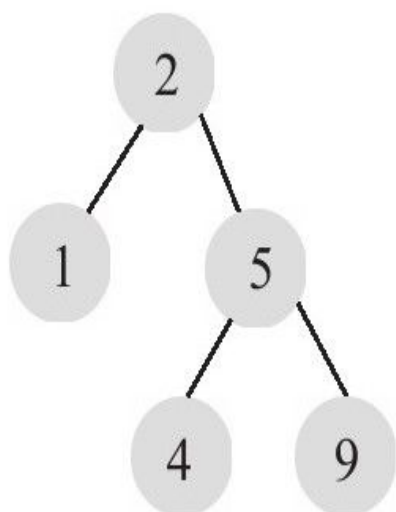
005-4000000



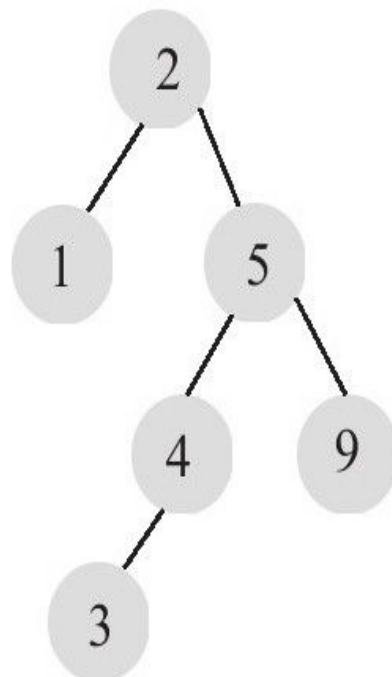
插入新值9



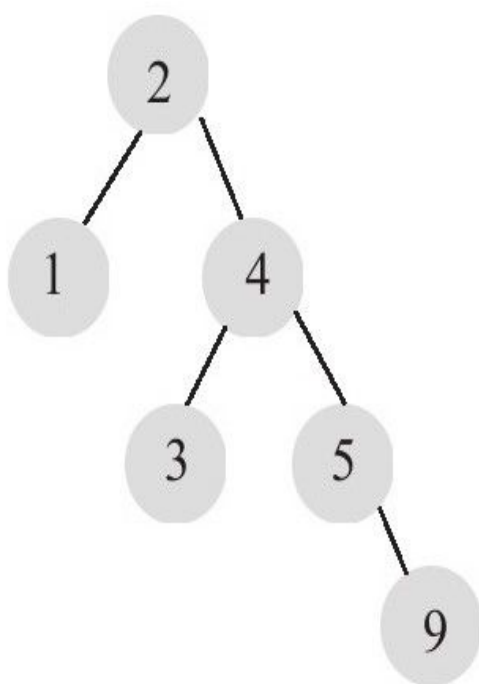
左旋以保证平衡



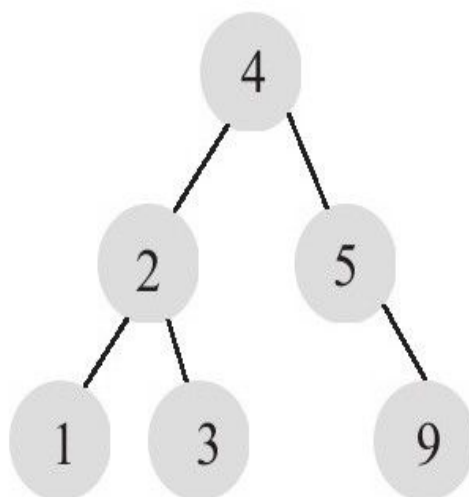
平衡二叉树



插入新键值3



右旋一次

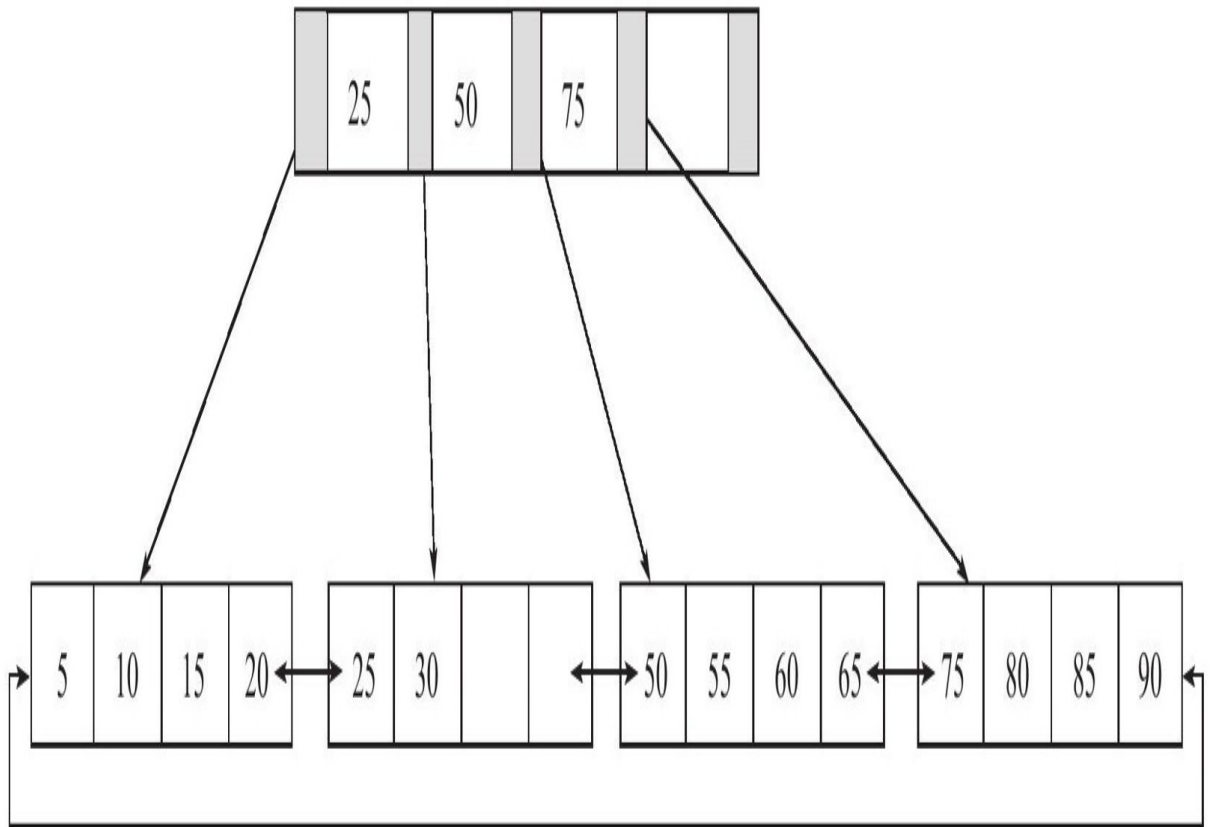


再左旋一次

5.3 B+

B+ ISAM B+ ISAM
MyISAM B

B+ B+ B+ B+
B+ 2 4 fan out 5 5-6



5-6 2 B+

5-6 5 10 15 20 25 30 50 55 60 65 75 80

85 90

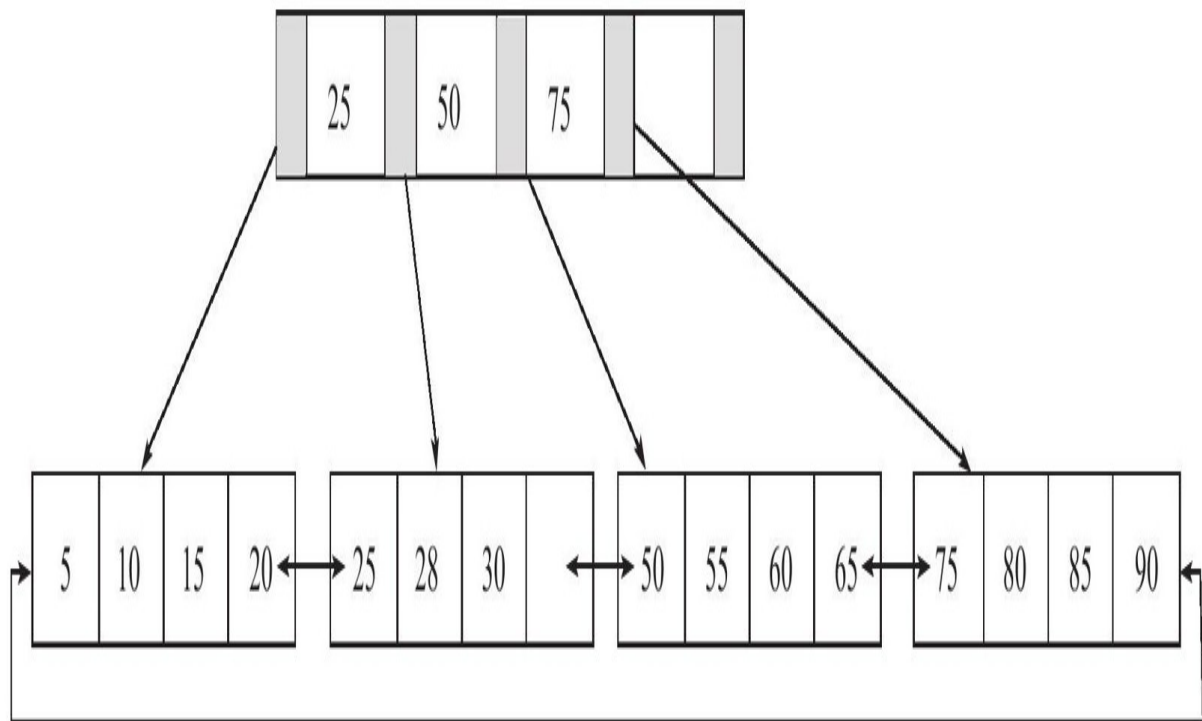
5.3.1 B+

B+ B+
5-1

表 5-1 B+ 树插入的 3 种情况

| Leaf Page 满 | Index Page 满 | 操作 |
|-------------|--------------|---|
| No | No | 直接将记录插入到叶子节点 |
| Yes | No | 1) 拆分 Leaf Page
2) 将中间的节点放入到 Index Page 中
3) 小于中间节点的记录放左边
4) 大于或等于中间节点的记录放右边 |
| Yes | Yes | 1) 拆分 Leaf Page
2) 小于中间节点的记录放左边
3) 大于或等于中间节点的记录放右边
4) 拆分 Index Page
5) 小于中间节点的记录放左边
6) 大于中间节点的记录放右边
7) 中间节点放入上一层 Index Page |

□□□□□□□□B+□□□□□□□□5-6□□□□B+□□□□□□28□□□□□□□□
 Leaf Page□Index Page□□□□□□□□□□□□□□□□5-7□

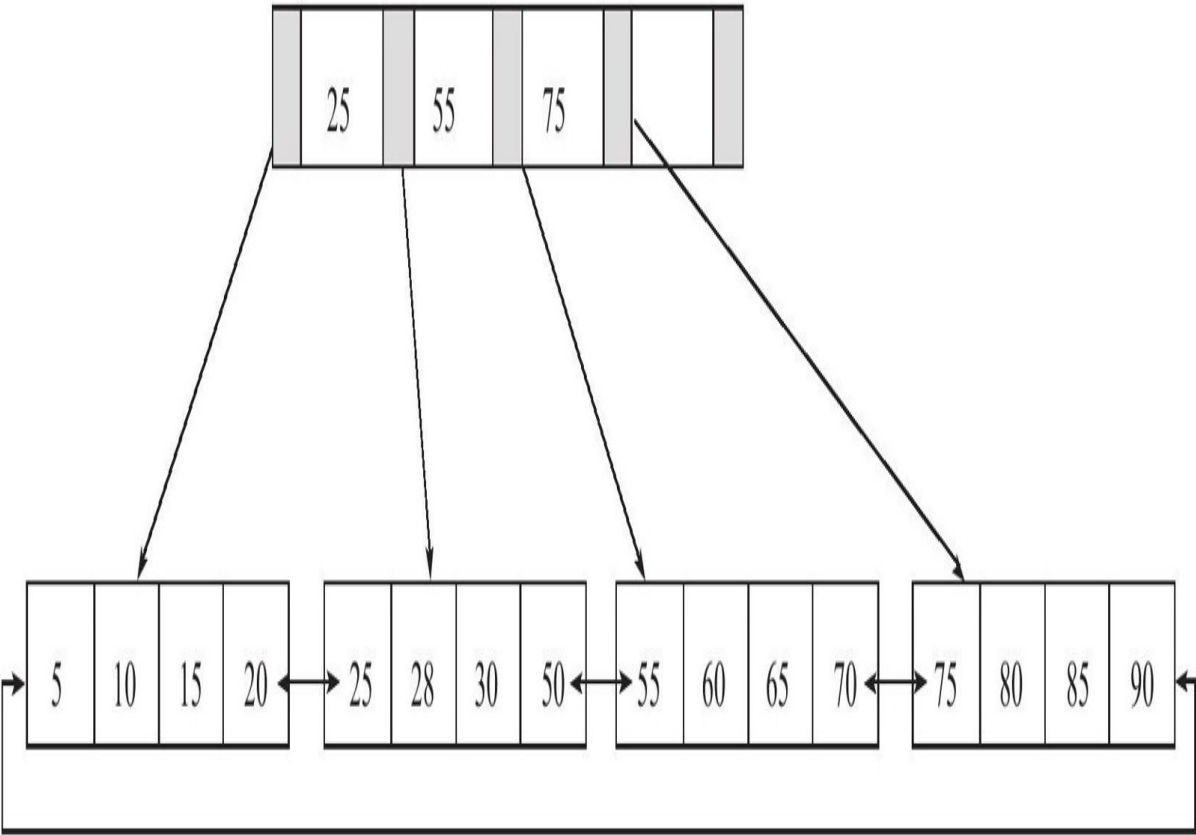


□ 5-7 □□□28

□□□□70□□□□□□□□Leaf Page□□□□□□Index Page□□□□□□5-1□□□□□□□□Leaf Page□□□□55□55□60□65□70□□□□□□□60□□□□□□□□5-8□

`split` B+ Rotation

Leaf Page
 B+
 5-7 70 B+ 5-10
 5-10



5-10 B+

5-10 B+ B+ 2

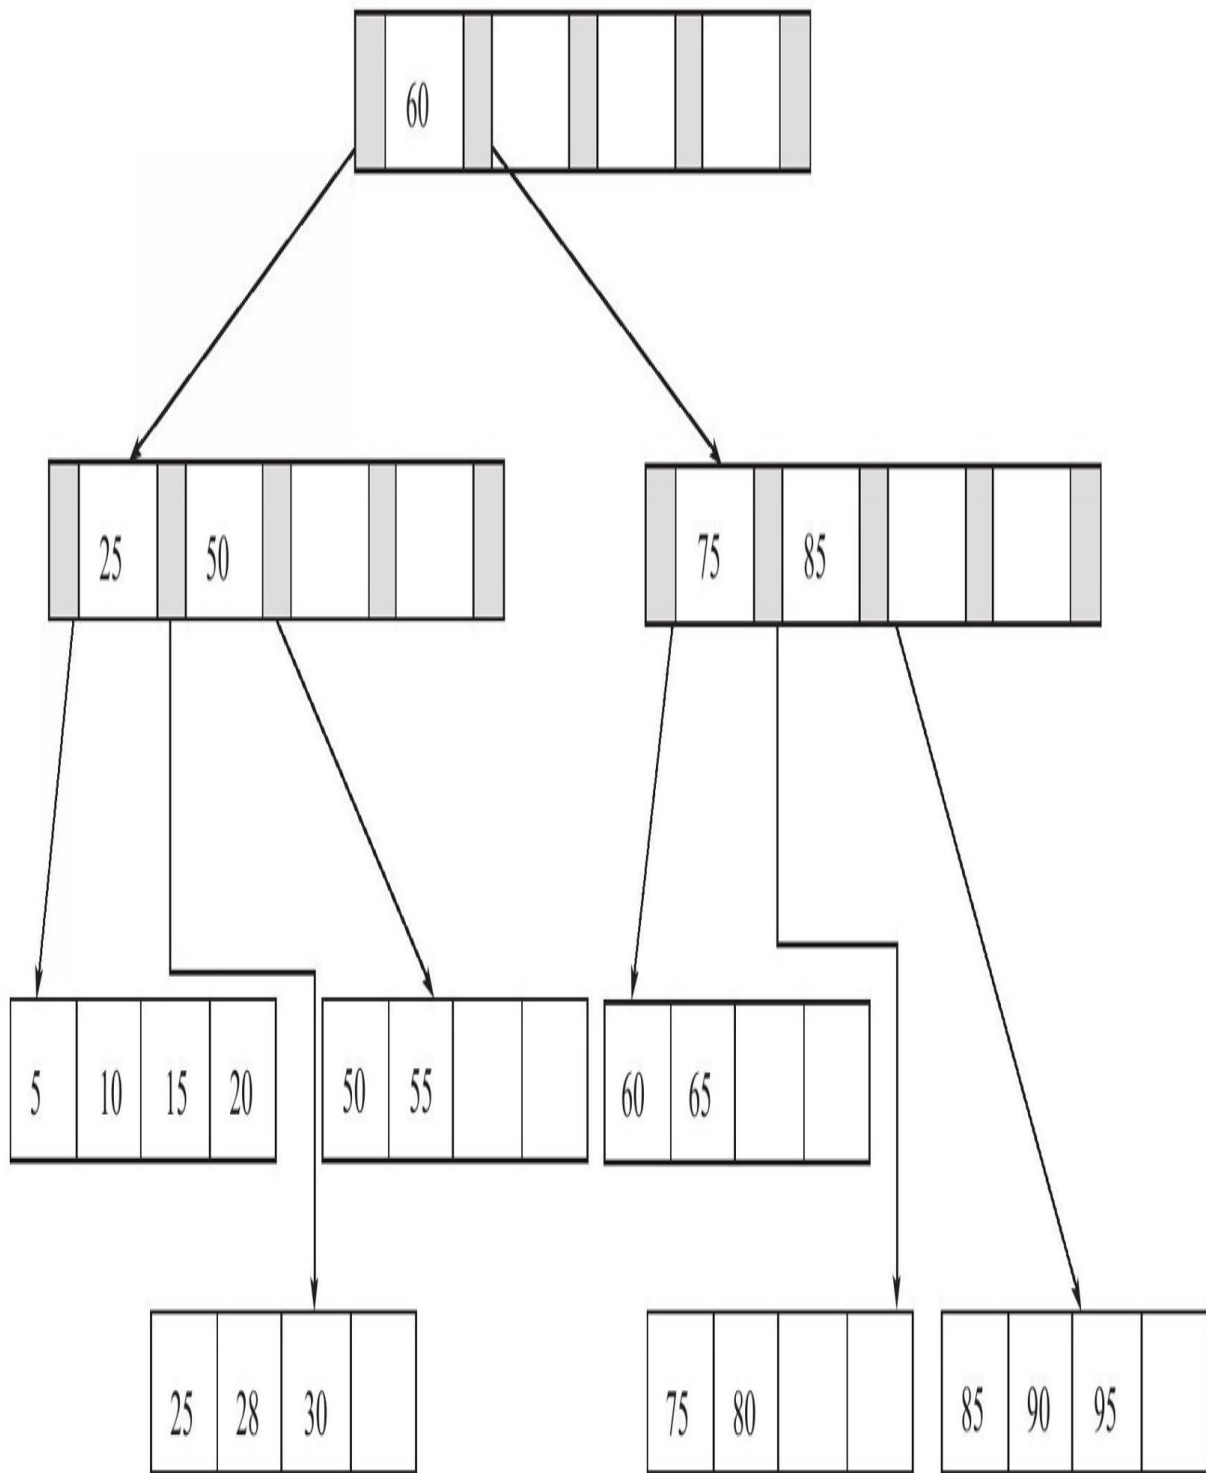
5.3.2 B+树删除

B+树的填充因子(fill factor)通常为50%~70%。B+树的删除操作分为三种情况：
1. 叶子节点小于填充因子
2. 中间节点小于填充因子
3. 叶子节点和中间节点都小于填充因子

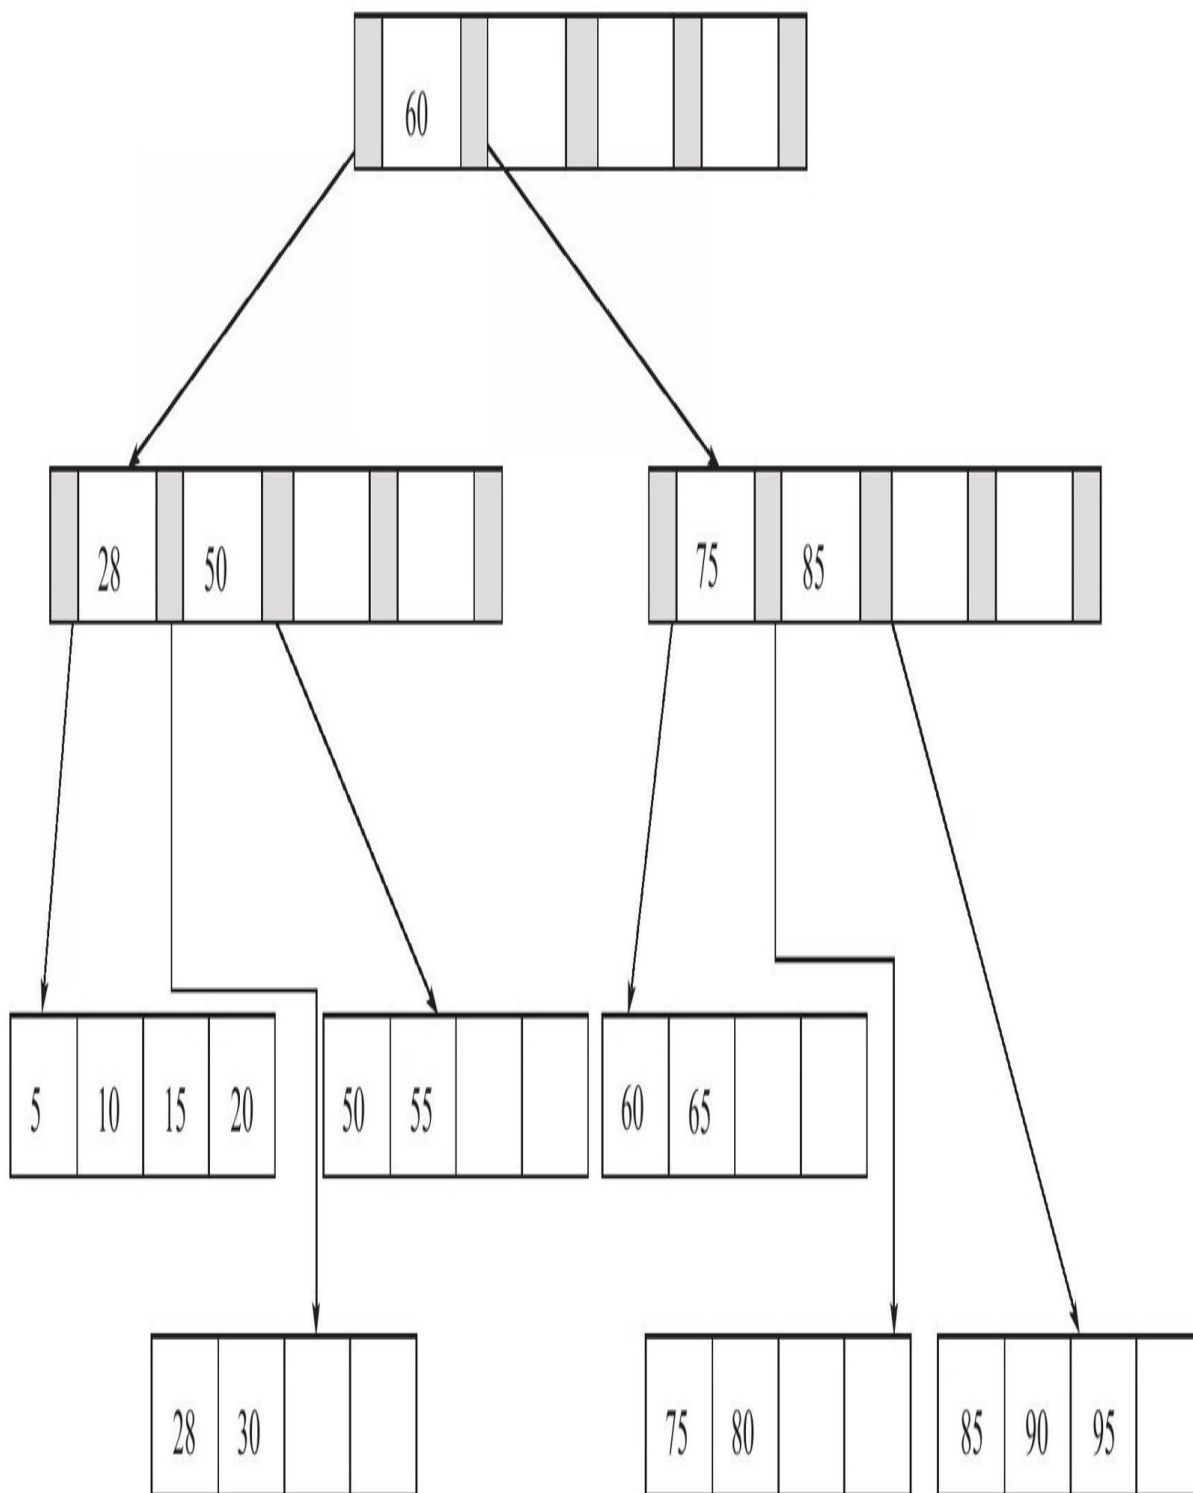
表 5-2 B+ 树删除操作的三种情况

| 叶子节点小于填充因子 | 中间节点小于填充因子 | 操作 |
|------------|------------|--|
| No | No | 直接将记录从叶子节点删除，如果该节点还是 Index Page 的节点，用该节点的右节点代替 |
| Yes | No | 合并叶子节点和它的兄弟节点，同时更新 Index Page |
| Yes | Yes | 1) 合并叶子节点和它的兄弟节点
2) 更新 Index Page
3) 合并 Index Page 和它的兄弟节点 |

图 5-9 B+ 树删除操作的三种情况
图 5-11 B+ 树删除操作的三种情况

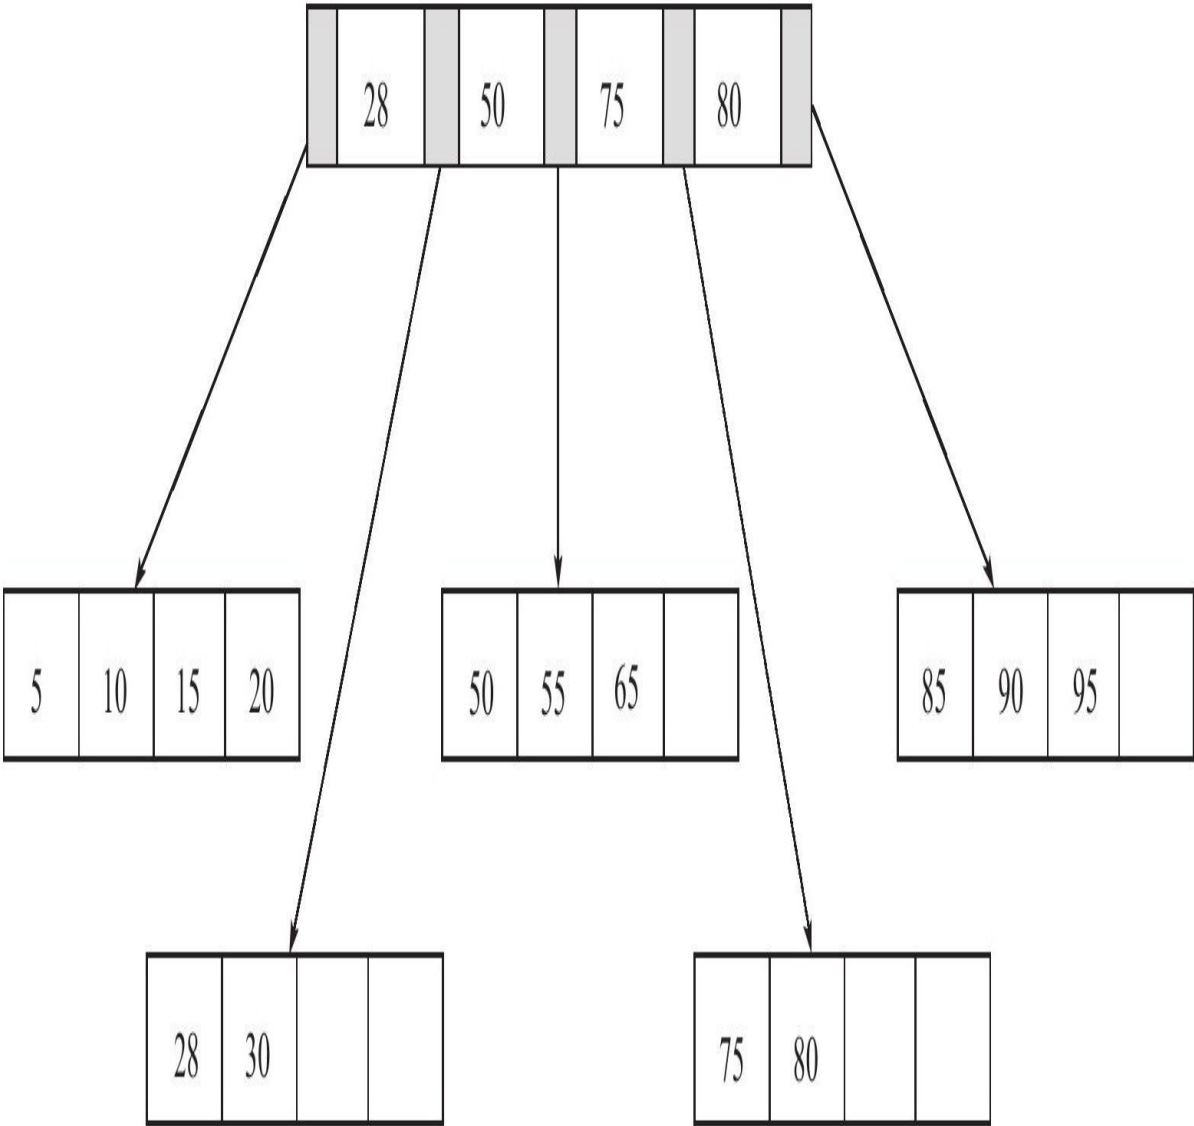


□□□□□□25□□□□□□□5-2□□□□□□□□□□□□□□□□IndexPage□□□□□□□□
□Leaf Page□□25□□□□□25□□□□□□□28□□□□Page Index□□□□□□□5-
12□



□ 5-12 □□□□25

60 Leaf Page 60 Fill Factor 50%
 Index Page Index Page
 5-13



5-13 60


```
INSERT INTO t SELECT 3,REPEAT('a',7000),-3;
```

```
INSERT INTO t SELECT 4,REPEAT('a',7000),-4;
```

py_innodb_page_info.py 7000
py_innodb_page_info

```
[root@nineyou0-43 data]#py_innodb_page_info.py-v mytest/t.ibd
```

```
page offset 00000000,page typeFile Space Header
```

```
page offset 00000001,page typeInsert Buffer Bitmap
```

```
page offset 00000002,page typeFile Segment inode
```

```
page offset 00000003,page typeB-tree Node,page level0001
```

```
page offset 00000004,page typeB-tree Node,page level0000
```

```
page offset 00000005,page typeB-tree Node,page level0000
```

```
page offset 00000006,page typeB-tree Node,page level0000
```

```
page offset 00000000,page typeFreshly Allocated Page
```

```
Total number of page:8:
```

```
Freshly Allocated Page:1
```

```
Insert Buffer Bitmap:1
```

```
File Space Header:1
```

```
B-tree Node:4
```

```
File Segment inode:1
```

page level0000
page level0001B+2B+
hexdump

```
0000c000 c2 33 62 95 00 00 00 03 ff ff ff ff ff ff ff|.3b.....|
```

```
0000c010 00 00 00 0a b6 8c ce 57 45 bf 00 00 00 00 00 00|......WE.....|
```

```
0000c020 00 00 00 00 00 f9 00 02 00 a2 80 05 00 00 00 00|......|
```

```
0000c030 00 9a 00 02 00 02 00 03 00 00 00 00 00 00 00 00|......|
```

0000c040 00 01 00 00 00 00 00 01 e2 00 00 00 f9 00 00|.....|
0000c050 00 02 00 f2 00 00 00 f9 00 00 02 00 32 01 00|.....2..|
0000c060 02 00 1b 69 6e 66 69 6d 75 6d 00 04 00 0b 00 00|...infimum.....|
0000c070 73 75 70 72 65 6d 75 6d 00 10 00 11 00 0e 80 00|supremum.....|
0000c080 00 01 00 00 00 04 00 00 00 19 00 0e 80 00 02|.....|
0000c090 00 00 00 05 00 00 00 21 ff d6 80 00 00 04 00 00|.....!.....|
0000c0a0 00 06 00 00 00 00 00 00 00 00 00 00 00 00|.....|
0000c0b0 00 00 00 00 00 00 00 00 00 00 00 00 00 00|.....|
0000c0c0 00 00 00 00 00 00 00 00 00 00 00 00 00 00|.....|
.....
0000fff0 00 00 00 00 00 70 00 63 73 d8 52 3a b6 8c ce 57|.....p.cs.R:...W|

Page Directory00 63Recorder Header0xc06369 6e 66 69 6d 75 6d
00infimum501 00 02 00 1bRecorder
Header481InnoDBPage
DirectoryinfimumRecorder Header
00 1b0xc063+1b=c07e80 00 00 01
1INT0x80 00 00 01
0x000180 00 00 0100 00 00 04
80 00 00 0280 00 00 04

5-14

mysql> EXPLAIN
mysql> SELECT * FROM Profile ORDER BY id LIMIT 10\G;

mysql> EXPLAIN
mysql> SELECT * FROM Profile ORDER BY id LIMIT 10\G; B+
mysql> EXPLAIN
mysql> SELECT * FROM Profile ORDER BY id LIMIT 10\G; EXPLAIN

mysql> EXPLAIN

- SELECT * FROM Profile ORDER BY id LIMIT 10\G;

*****1.row*****

id:1

select_type:SIMPLE

table:Profile

type:index

possible_keys:NULL

key:PRIMARY

key_len:4

ref:NULL

rows:10

Extra:

1 row in set(0.00 sec)

mysql> EXPLAIN
mysql> SELECT * FROM Profile ORDER BY id LIMIT 10\G; filesort

mysql> EXPLAIN
mysql> SELECT * FROM Profile ORDER BY id LIMIT 10\G; range query

mysql> EXPLAIN

- SELECT * FROM Profile

- WHERE id < 10 AND id > 10000\G;

*****1.row*****

```
id:1

select_type:SIMPLE

table:Profile

type:range

possible_keys:PRIMARY

key:PRIMARY

key_len:4

ref:NULL

rows:14868

Extra:Using where

1 row in set (0.01 sec)
```

EXPLAIN MySQL execute plan rows
rows
SQL 9946

```
mysql> SELECT COUNT(*) from Profile

- WHERE id<10 AND id<10000;

*****1.row*****

COUNT(1):9946

1 row in set (0.00 sec)
```

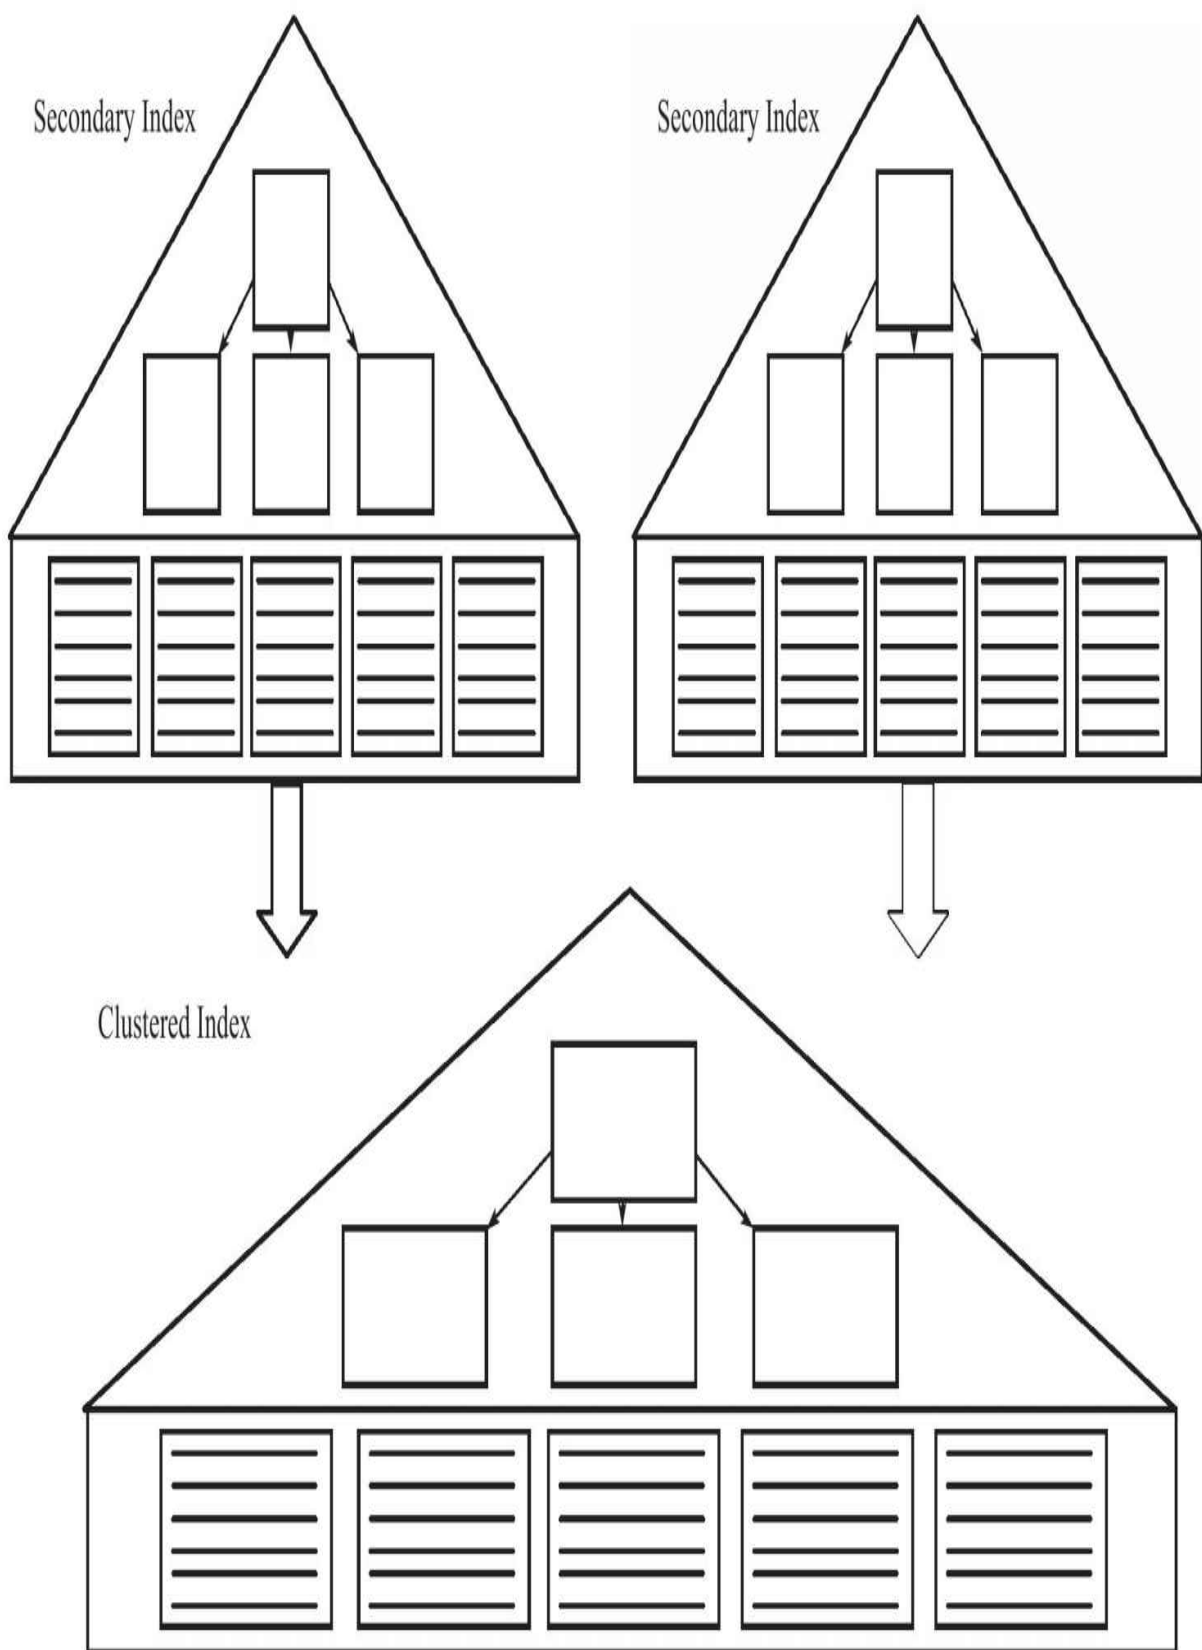
[1] non-clustered index

5.4.2 〇〇〇〇

```

Secondary Index
bookmark
InnoDB
InnoDB
InnoDB
5-15
InnoDB

```



5-15

[illegible]

```

Microsoft SQL Server
MySQL
MyISAM
NOT NULL
Row Identifiedr
RID
“ ”

```

Microsoft SQL Server DBA Microsoft SQL Server

OLAP On-Line Analytical Processing
B+
read ahead
InnoDB MyISAM
It all depends

```

InnoDB
t
t
c
c

```

```
mysql> ALTER TABLE t
```

```
- ADD c INT NOT NULL;
```

Query OK, 4 rows affected (0.24 sec)

Records:4 Duplicates:0 Warnings:0

```
mysql> UPDATE t SET c=0-a;
```

Query OK, 4 rows affected (0.04 sec)

Rows matched:4 Changed:4 Warnings:0

mysql>ALTER TABLE t ADDKEY idx_c(c);

Query OK,4 rows affected(0.28 sec)

Records:4 Duplicates:0 Warnings:0

mysql>SHOW INDEX FROM t\G;

*****1.row*****

Table:t

Non_unique:0

Key_name:PRIMARY

Seq_in_index:1

Column_name:a

Collation:A

Cardinality:2

Sub_part:NULL

Packed:NULL

Null:

Index_type:BTREE

Comment:

*****2.row*****

Table:t

Non_unique:1

Key_name:idx_c

Seq_in_index:1

Column_name:c

Collation:A

Cardinality:2

Sub_part:NULL

Packed:NULL

Null:

Index_type:BTREE

Comment:

2 rows in set 0.00 sec

mysql>select a,c from t;

+---+-----+

|a|c|

+---+-----+

|4|-4|

|3|-3|

|2|-2|

|1|-1|

+---+-----+

4 rows in set 0.00 sec

py_innodb_page_info

[root@nineyou0-43 mytest]#py_innodb_page_info.py -v t.ibd

page offset 00000000,page typeFile Space Header

page offset 00000001,page typeInsert Buffer Bitmap

page offset 00000002,page typeFile Segment inode

page offset 00000003,page typeB-tree Node,page level0001

page offset 00000004,page typeB-tree Node,page level0000

page offset 00000005,page typeB-tree Node,page level0000

page offset 00000006,page typeB-tree Node,page level0000

page offset 00000007,page typeB-tree Node,page level0000

page offset 00000000,page typeFreshly Allocated Page

Total number of page:9:

Freshly Allocated Page:1

Insert Buffer Bitmap:1

File Space Header:1

B-tree Node:5

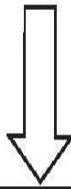
File Segment inode:1

page offset4
hexdump

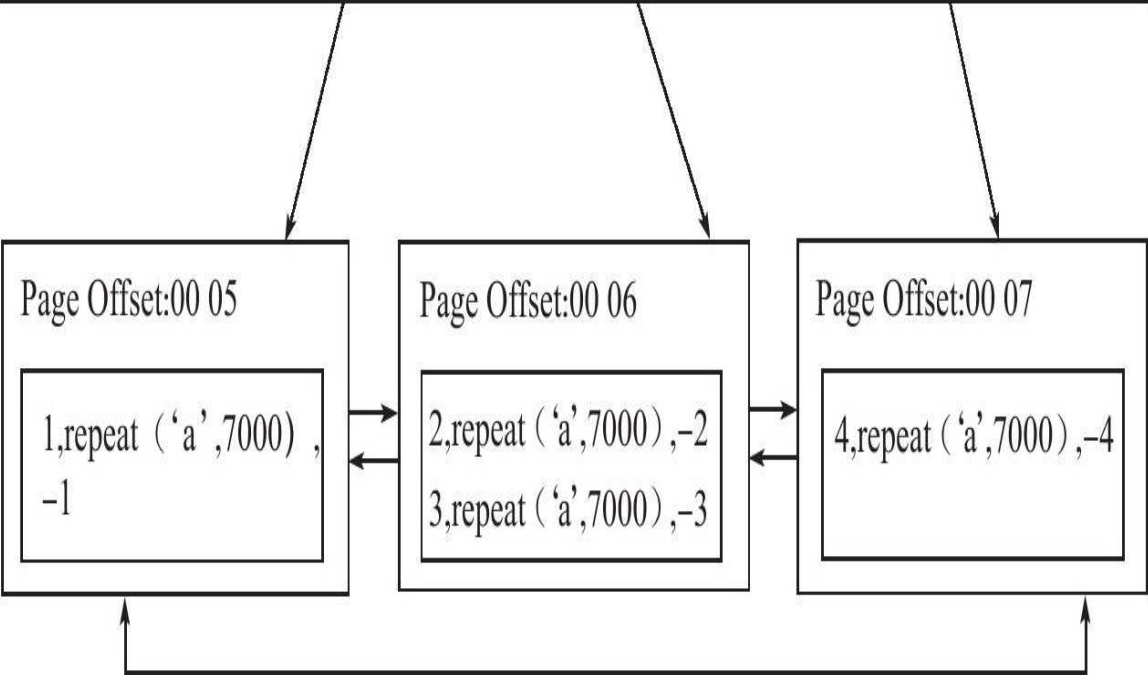
```
00010000 b9 aa 8e d0 00 00 00 04 ff ff ff ff ff ff ff|.....|
00010010 00 00 00 0a ec ea 4e 27 45 bf 00 00 00 00 00 00|.....N'E.....|
00010020 00 00 00 00 01 02 00 02 00 ac 80 06 00 00 00 00|.....|
00010030 00 a4 00 01 00 03 00 04 00 00 00 00 00 52 d4 8b|.....R..|
00010040 00 00 00 00 00 00 00 00 01 f2 00 00 01 02 00 00|.....|
00010050 00 02 02 72 00 00 01 02 00 00 00 02 01 b2 01 00|...r.....|
00010060 02 00 41 69 6e 66 69 6d 75 6d 00 05 00 0b 00 00|..Ainfimum.....|
00010070 73 75 70 72 65 6d 75 6d 00 00 10 ff f3 7f ff ff|supremum.....|
00010080 ff 80 00 00 01 00 00 18 ff f3 7f ff ff fe 80 00|.....|
00010090 00 02 00 00 20 ff f3 7f ff ff fd 80 00 00 03 00|.....|
000100a0 00 28 ff f3 7f ff ff fc 80 00 00 04 00 00 00 00|.□.....|
.....
00013ff0 00 00 00 00 00 70 00 63 f3 46 77 f2 ec ea 4e 27|.....p.C.Fw...N'|
```

4c45-16

| | | | |
|--|--|--|--|
| 辅助索引idx_c | | | |
| Page Offset:00 04 | | | |
| Key:7f ff ff ff
Pointer:80 00 00 01 | Key:7f ff ff fe
Pointer:80 00 00 02 | Key:7f ff ff fd
Pointer:80 00 00 03 | Key:7f ff ff fc
Pointer:80 00 00 04 |



| | | | |
|----------------------------------|----------------------------------|----------------------------------|--|
| 聚集索引 Page Offset:00 03 | | | |
| Key:80 00 00 01
Pointer:00 05 | Key:80 00 00 02
Pointer:00 06 | Key:80 00 00 04
Pointer:00 07 | |



5-16

5-16 t idx_c c
-1 7f ff ff ff 7 0111
0 1 -1

5.4.3 B+□□□□□□

[illegible][illegible]

1 2 3 4 5 6 7 8 9

10 5.3.1
5 split record

P1 1 2 3 4

P2□5□6□7□8□9□10

P1P2

InnoDBPage Header

☐PAGE_LAST_INSERT

☐PAGE_DIRECTION

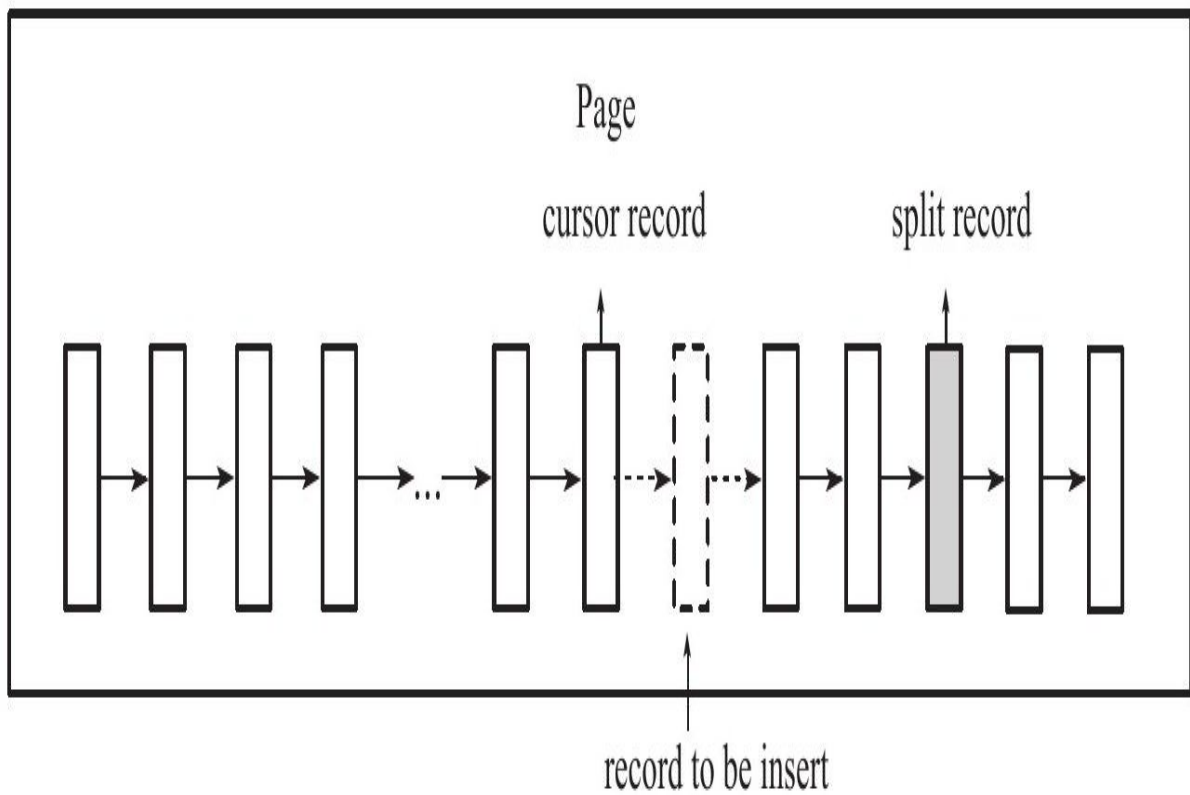
☐PAGE_N_DIRECTION

```

InnoDB
5
cursor
InnoDB
3

```

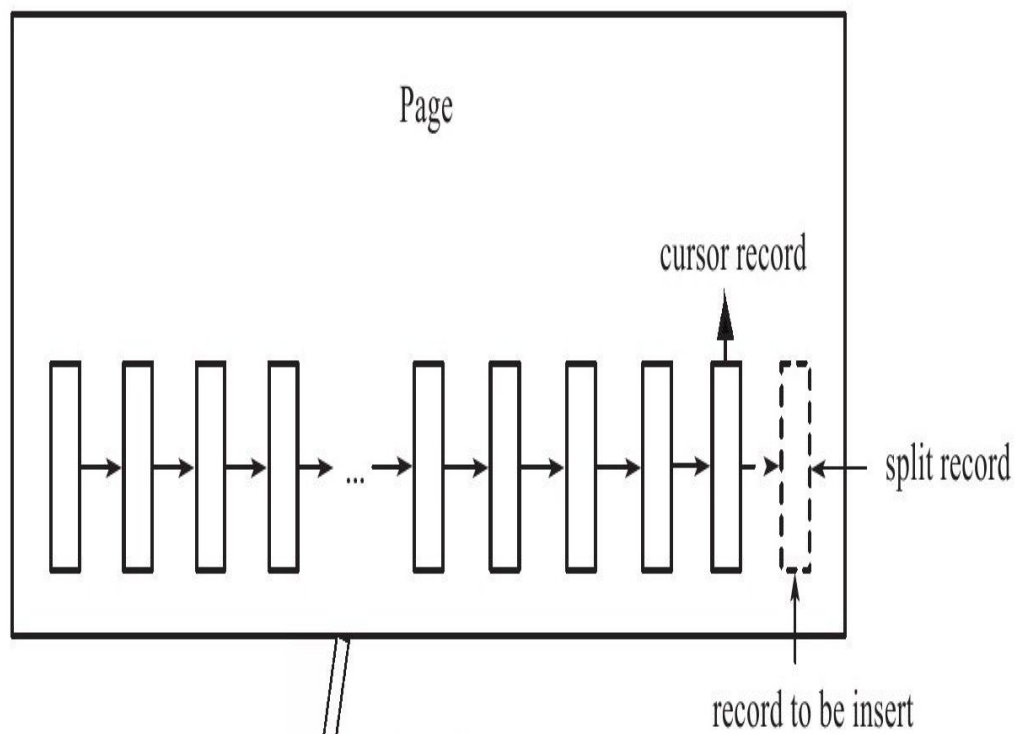
35-17



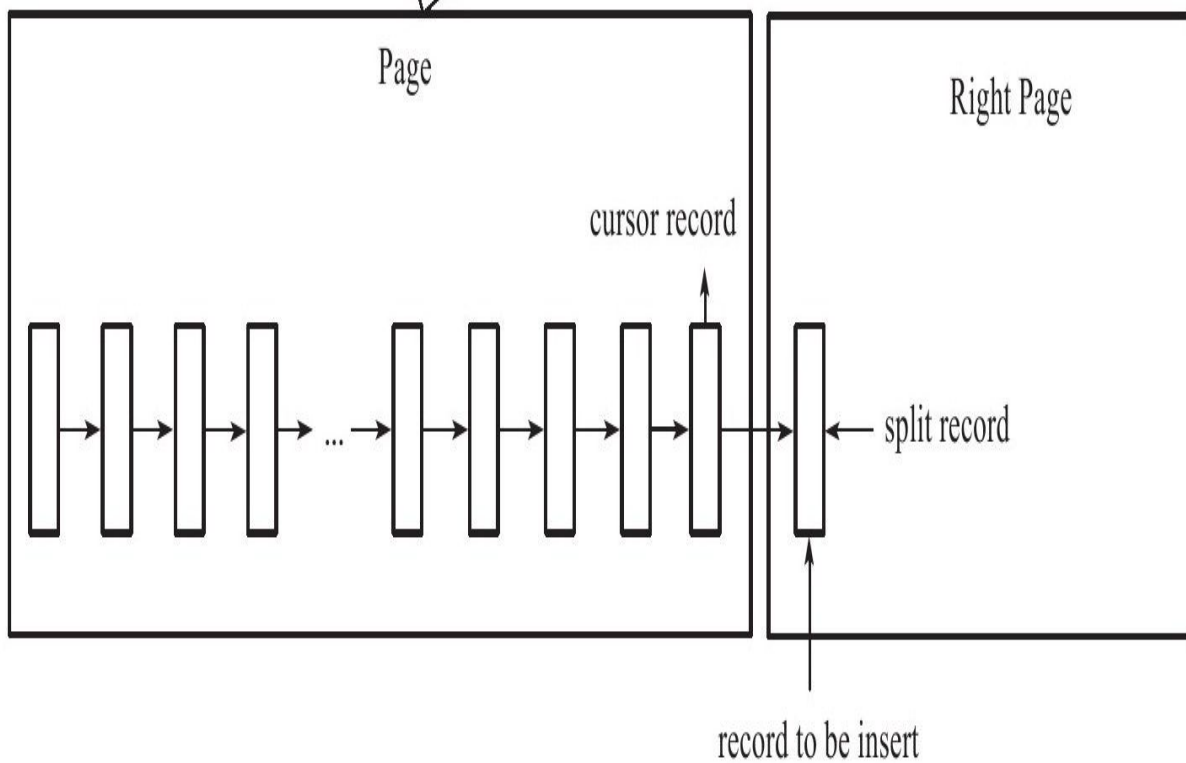
□ 5-17 □□□□□□□□

□5-17□□□□□□□□□□□□□□3□□□□split record□□□□□□□□□□□□□□

5-18□□□□□□



After Split



□ 5-19 □□□□□□□□

5.4.4 B+

1.

ALTER TABLE CREATE/DROP INDEX ALTER TABLE

```
ALTER TABLE tbl_name  
  
|ADD{INDEX|KEY}[index_name]  
  
[index_type](index_col_name,...)[index_option]...  
  
ALTER TABLE tbl_name  
  
DROP PRIMARY KEY  
  
|DROP{INDEX|KEY}index_name
```

CREATE/DROP INDEX

```
CREATE[UNIQUE]INDEX index_name  
  
[index_type]  
  
ON tbl_name(index_col_name,...)  
  
DROP INDEX index_name ON tbl_name
```

t b
varchar8000100

```
mysql> ALTER TABLE t  
  
->ADD KEY idx_b(b(100));  
  
Query OK, 4 rows affected (0.32 sec)  
  
Records: 4 Duplicates: 0 Warnings: 0
```

mysql> SHOW INDEX FROM t
mysql> ALTER TABLE t
ADD KEY idx_a_c(a,c);

mysql> ALTER TABLE t

-> ADD KEY idx_a_c(a,c);

Query OK, 4 rows affected (0.31 sec)

Records: 4 Duplicates: 0 Warnings: 0

mysql> SHOW INDEX FROM t\G;

*****1.row*****

Table: t

Non_unique: 0

Key_name: PRIMARY

Seq_in_index: 1

Column_name: a

Collation: A

Cardinality: 2

Sub_part: NULL

Packed: NULL

Null:

Index_type: BTREE

Comment:

*****2.row*****

Table: t

Non_unique: 1

Key_name: idx_b

Seq_in_index: 1

Column_name: b

Collation: A

Cardinality: 2

Sub_part: 100

Packed:NULL

Null:YES

Index_type:BTREE

Comment:

*****3.row*****

Table:t

Non_unique:1

Key_name:idx_a_c

Seq_in_index:1

Column_name:a

Collation:A

Cardinality:2

Sub_part:NULL

Packed:NULL

Null:

Index_type:BTREE

Comment:

*****4.row*****

Table:t

Non_unique:1

Key_name:idx_a_c

Seq_in_index:2

Column_name:c

Collation:A

Cardinality:2

Sub_part:NULL

Packed:NULL

Null:

Index_type:BTREE

Comment:

*****5.row*****

Table:t

Non_unique:1

Key_name:idx_c

Seq_in_index:1

Column_name:c

Collation:A

Cardinality:2

Sub_part:NULL

Packed:NULL

Null:

Index_type:BTREE

Comment:

5 rows in set (0.00 sec)

SHOW INDEX FROM t 4 c
b 100 a c SHOW
INDEX

Table

Non_unique primary key 0

Key_name DROP INDEX

Seq_in_index idx_a_c

Column_name

Collation A NULL B+ A
Heap Hash NULL Hash Hash
Hash Hash

Cardinality Cardinality
1

Collation:A
Cardinality:4
Sub_part:NULL
Packed:NULL
Null:
Index_type:BTREE
Comment:

.....

CardinalityANALYZE
TABLEMySQL
Cardinality

mysqlshow index from Profile\G;

*****1.row*****

Table:Profile

Non_unique:0

Key_name:UserName

Seq_in_index:1

Column_name:username

Collation:A

Cardinality:NULL

Sub_part:NULL

Packed:NULL

Null:

Index_type:BTREE

Comment:

CardinalityNULL
EXPLAIN

ANALYZE TABLE 语句在 MySQL 5.5 版本之前是不支持的，在 MySQL 5.5 版本之后，ANALYZE TABLE 语句被重新引入，用于分析表的索引。

2. Fast Index Creation

MySQL 5.5 版本之前，创建索引是一个非常慢的过程，因为 MySQL 需要扫描整个表来创建索引。在 MySQL 5.5 版本之后，引入了 Fast Index Creation 功能，大大加快了索引的创建速度。

在 MySQL 5.5 版本之后，可以使用 ALTER TABLE 语句来创建索引。

例如，以下语句用于在表 mytable 上创建索引：

```
ALTER TABLE mytable ADD INDEX (column_name);
```

在 MySQL 5.5 版本之前，创建索引需要使用 CREATE INDEX 语句。

在 MySQL 5.5 版本之前，创建索引是一个非常慢的过程，因为 MySQL 需要扫描整个表来创建索引。在 MySQL 5.5 版本之后，引入了 Fast Index Creation 功能，大大加快了索引的创建速度。Microsoft SQL Server 和 Oracle 数据库也有类似的功能，但 MySQL 的 Fast Index Creation 功能更加强大。

InnoDB 引擎支持 Fast Index Creation 功能。在 InnoDB 1.0.x 版本之前，Fast Index Creation 功能是不支持的。在 InnoDB 1.0.x 版本之后，Fast Index Creation 功能被引入，大大加快了索引的创建速度。

InnoDB 引擎支持 Fast Index Creation 功能。在 InnoDB 1.0.x 版本之前，Fast Index Creation 功能是不支持的。在 InnoDB 1.0.x 版本之后，Fast Index Creation 功能被引入，大大加快了索引的创建速度。MySQL 5.5 版本之后，Fast Index Creation 功能被引入，大大加快了索引的创建速度。

在 MySQL 5.5 版本之前，创建索引是一个非常慢的过程，因为 MySQL 需要扫描整个表来创建索引。在 MySQL 5.5 版本之后，引入了 Fast Index Creation 功能，大大加快了索引的创建速度。

Fast Index Creation 功能在 MySQL 5.5 版本之后被引入，大大加快了索引的创建速度。在 MySQL 5.5 版本之前，创建索引是一个非常慢的过程，因为 MySQL 需要扫描整个表来创建索引。

3. Online Schema Change

Online Schema Change [OSC](#) [Facebook](#)
DDL [Facebook](#) MySQL “ ”
MySQL DDL

Facebook PHP OSC InnoDB OSC
Facebook Vamsi Ponnekanti OSC
The openarkkit toolkit oak-online-alter-table OSC

☐ init

☐ createCopyTable

☐ alterCopyTable ALTER TABLE

☐ createDeltasTable deltas
DML createDeltasTable

☐ createTriggers INSERT UPDATE DELETE
deltas

☐ startSnapshotXact OSC

☐ selectTableIntoOutfile
chunked copy
500 000

☐ dropNCIndexs

☐ loadCopyTable

☐ replayChanges OSC DML
deltas

☐ recreateNCIndexes

☐ replayChanges DML

❑ swapTables 函数在 2 个数据库之间交换表

OSC 的 PHP 版本 2200 支持 MySQL InnoDB DBA 工具对 InnoDB 数据库

OSC 的 PHP 版本 2200 支持 SET sql_bin_log=0 关闭主从复制

4. Online DDL

FIC 的 InnoDB 数据库支持在线 DDL 操作，MySQL 5.6 支持 Online DDL 操作，INSERT、UPDATE、DELETE 等 DML 操作在 MySQL 5.6 中支持

DDL 操作在“DDL”菜单中

❑ 数据库

❑ 表

❑ 索引

❑ 视图

❑ ALTER TABLE 语句

```
ALTER TABLE tbl_name
```

```
[ADD{INDEX|KEY}{index_name}
```

```
[index_type](index_col_name,...)[index_option]...
```

```
ALGORITHM={DEFAULT|INPLACE|COPY}
```

```
LOCK={DEFAULT|NONE|SHARED|EXCLUSIVE}
```

ALGORITHM은 테이블을 복사할 때 MySQL 5.1 이전 버전에서는
INPLACE로, 5.1 이후 버전에서는 DEFAULT로
old_alter_table은 INPLACE로 COPY로, OFF로
INPLACE로

```
mysql> SELECT @@version\G;
```

```
*****1.row*****
```

```
@@version:5.6.6-m9
```

```
1 row in set (0.00 sec)
```

```
mysql> SHOW VARIABLES LIKE 'old_alter_table'\G;
```

```
*****1.row*****Variable_name:old_alter_table
```

```
Value:OFF
```

```
1 row in set (0.00 sec)
```

LOCK은 테이블을 잠그는 방식

1 NONE

테이블을 잠그지 않고 읽기 전용으로
테이블을 잠그지 않고 쓰기 전용으로

2 SHARE

테이블을 FIC로 잠그고 쓰기 전용으로 S로 잠그고 쓰기 전용으로
테이블을 잠그고 쓰기 전용으로 SHARE로 잠그고 쓰기 전용으로

3 EXCLUSIVE

EXCLUSIVE로 잠그고 쓰기 전용으로 X로 잠그고 쓰기 전용으로
테이블을 잠그고 쓰기 전용으로 COPY로 잠그고 쓰기 전용으로

4 DEFAULT

DEFAULT NONE SHARE
EXCLUSIVE DEFAULT
DDL

InnoDB Online DDL INSERT
UPDATE DELETE DML
innodb_online_alter_log_max_size
128MB
innodb_online_alter_log_max_size

Error:1799SQLSTATE:HY000(ER_INNODB_ONLINE_LOG_TOO_BIG)

Message:Creating index'idx_aaa'required more than'innodb_online_alter_log_max_size'bytes of modification log.Please try again.

innodb_online_alter_log_max_size
ALTER TABLE SHARE
DML

Online DDL
SQL

5.5 Cardinality

5.5.1 B+ Cardinality

B+ 트리의 Cardinality는 트리의 높이와 노드의 수에 따라 달라진다. 트리의 높이가 높을수록 Cardinality는 증가한다. B+ 트리의 Cardinality는 트리의 높이와 노드의 수에 따라 달라진다. 트리의 높이가 높을수록 Cardinality는 증가한다.

```
SELECT * FROM student WHERE sex='M'
```

SQL에서 WHERE 절을 사용하여 데이터를 필터링할 때, 트리의 Cardinality는 트리의 높이와 노드의 수에 따라 달라진다. 트리의 높이가 높을수록 Cardinality는 증가한다. B+ 트리의 Cardinality는 트리의 높이와 노드의 수에 따라 달라진다. 트리의 높이가 높을수록 Cardinality는 증가한다.

SHOW INDEX 명령을 사용하여 트리의 Cardinality를 확인할 수 있다. Cardinality는 트리의 높이와 노드의 수에 따라 달라진다. 트리의 높이가 높을수록 Cardinality는 증가한다. B+ 트리의 Cardinality는 트리의 높이와 노드의 수에 따라 달라진다. 트리의 높이가 높을수록 Cardinality는 증가한다.

```
SELECT * FROM member WHERE usernick='David'
```

member 테이블의 500개의 usernick 값 중에서 David라는 값을 찾는 데에 B+ 트리의 Cardinality는 트리의 높이와 노드의 수에 따라 달라진다. 트리의 높이가 높을수록 Cardinality는 증가한다. B+ 트리의 Cardinality는 트리의 높이와 노드의 수에 따라 달라진다. 트리의 높이가 높을수록 Cardinality는 증가한다.

```
mysql> EXPLAIN SELECT * FROM member
```

```
- WHERE usernick='David'\G;
```

```
*****1.row*****
```

```
id:1
```

select_type:SIMPLE

table:member

type:const

possible_keys:username

key:username

key_len:62

ref:const

rows:1

Extra:

1 row in set(0.00 sec)

username SQL

5.5.2 InnoDB Cardinality

MySQL Cardinality Cardinality Cardinality Cardinality
MySQL Cardinality
MySQL B+ Cardinality

Cardinality
50G Cardinality
Cardinality Sample

InnoDB Cardinality INSERT
UPDATE INSERT UPDATE Cardinality
InnoDB Cardinality

1/16

stat_modified_counter 2 000 000 000

Cardinality 1/16
Cardinality
InnoDB
stat_modified_counter
stat_modified_counter 2 000 000 000 Cardinality

InnoDB Cardinality
InnoDB 8 Leaf Page

B+ A

B+ 8 P1 P2 ... P8

Cardinality Cardinality = $P_1 + P_2 + \dots + P_8$
 $\times A/8$

InnoDB Cardinality 8
 Cardinality 8
 Cardinality Cardinality

```
SHOW INDEX FROM OrderDetails
```

SQL MySQL Cardinality 5-20

| | Table | Non_unique | Key_name | Seq_in_index | Column_name | Collation | Cardinality | Sub_part | Packed | Null | Index_type | Comment | Index_comment |
|---|--------------|------------|-----------------------|--------------|-------------|-----------|-------------|----------|--------|------|------------|---------|---------------|
| ▶ | orderdetails | 0 | PRIMARY | 1 | OrderID | A | 2032 | NULL | NULL | | BTREE | | |
| | orderdetails | 0 | PRIMARY | 2 | ProductID | A | 2032 | NULL | NULL | | BTREE | | |
| | orderdetails | 1 | OrderID | 1 | OrderID | A | 2032 | NULL | NULL | | BTREE | | |
| | orderdetails | 1 | OrdersOrder_Details | 1 | OrderID | A | 2032 | NULL | NULL | | BTREE | | |
| | orderdetails | 1 | ProductID | 1 | ProductID | A | 156 | NULL | NULL | | BTREE | | |
| | orderdetails | 1 | ProductsOrder_Details | 1 | ProductID | A | 156 | NULL | NULL | | BTREE | | |

5-20 SHOW INDEX FROM OrderDetails

INSERT UPDATE OrderDetails
 SHOW INDEX FROM Cardinality 5-21

| | Table | Non_unique | Key_name | Seq_in_index | Column_name | Collation | Cardinality | Sub_part | Packed | Null | Index_type | Comment | Index_comment |
|---|--------------|------------|-----------------------|--------------|-------------|-----------|-------------|----------|--------|------|------------|---------|---------------|
| ▶ | orderdetails | 0 | PRIMARY | 1 | OrderID | A | 2192 | NULL | NULL | | BTREE | | |
| | orderdetails | 0 | PRIMARY | 2 | ProductID | A | 2192 | NULL | NULL | | BTREE | | |
| | orderdetails | 1 | OrderID | 1 | OrderID | A | 2192 | NULL | NULL | | BTREE | | |
| | orderdetails | 1 | OrdersOrder_Details | 1 | OrderID | A | 2192 | NULL | NULL | | BTREE | | |
| | orderdetails | 1 | ProductID | 1 | ProductID | A | 168 | NULL | NULL | | BTREE | | |
| | orderdetails | 1 | ProductsOrder_Details | 1 | ProductID | A | 168 | NULL | NULL | | BTREE | | |

5-21 SHOW INDEX FROM OrderDetails

SHOW INDEX FROM OrderDetails
 Cardinality OrderDetails
 Cardinality 8
 Cardinality InnoDB Bug

Cardinality
 8
 Cardinality

InnoDB 1.2 innodb_stats_sample_pages
 Cardinality 8 innodb_stats_method
 NULL nulls_equal NULL
 nulls_unequal nulls_ignored NULL
 NULL NULL 1 2 3 3 3
 innodb_stats_method Cardinality 4
 innodb_stats_method nulls_unequal Cardinality 5
 innodb_stats_method nulls_ignored Cardinality 3

SQL ANALYZE TABLE SHOW TABLE STATUS SHOW
 INDEX INFORMATION_SCHEMA TABLES STATISTICS
 InnoDB Cardinality
 Cardinality

InnoDB1.2 Cardinality 5-3

表 5-3 InnoDB 1.2 新增参数

| 参数 | 说明 |
|--------------------------------------|---|
| innodb_stats_persistent | <p>是否将命令 ANALYZE TABLE 计算得到的 Cardinality 值存放到磁盘上。若是，则这样做的好处是可以减少重新计算每个索引的 Cardinality 值，例如当 MySQL 数据库重启时。此外，用户也可以通过命令 CREATE TABLE 和 ALTER TABLE 的选项 STATS_PERSISTENT 来对每张表进行控制。</p> <p>默认值：OFF</p> |
| innodb_stats_on_metadata | <p>当通过命令 SHOW TABLE STATUS、SHOW INDEX 及访问 INFORMATION_SCHEMA 架构下的表 TABLES 和 STATISTICS 时，是否需要重新计算索引的 Cardinality 值。</p> <p>默认值：OFF</p> |
| innodb_stats_persistent_sample_pages | <p>若参数 innodb_stats_persistent 设置为 ON，该参数表示 ANALYZE TABLE 更新 Cardinality 值时每次采样页的数量。</p> <p>默认值：20</p> |
| innodb_stats_transient_sample_pages | <p>该参数用来取代之前版本的参数 innodb_stats_sample_pages，表示每次采样页的数量。</p> <p>默认值为：8</p> |

5.6 B+樹索引

5.6.1 為什麼需要B+樹索引

為什麼需要B+樹索引？
B+樹索引是資料庫中一種非常重要的索引結構。它與B樹索引類似，但有一些關鍵的區別。DBA（資料庫管理員）需要理解這些區別，因為它們會影響資料庫的性能和維護。Think Different。

1. 索引的類型：
OLTP（Online Transaction Processing）和OLAP（Online Analytical Processing）是兩種不同的資料庫工作負載。OLTP通常涉及大量的寫入操作，而OLAP則涉及大量的讀取操作。B+樹索引在OLTP環境中表現良好，但在OLAP環境中可能不是最佳選擇。B+樹索引的缺點是它在讀取大量數據時效率較低，因為它需要訪問多個索引塊。

OLAP（Online Analytical Processing）是另一種資料庫工作負載。OLAP通常涉及大量的讀取操作，而OLTP則涉及大量的寫入操作。OLAP索引的缺點是它在寫入操作時效率較低，因為它需要更新大量的索引塊。OLAP索引的優點是它在讀取大量數據時效率較高，因為它只需要訪問一個索引塊。Hash Join（哈希連接）是OLAP中常用的一種連接操作。DBA（資料庫管理員）需要理解這些區別，因為它們會影響資料庫的性能和維護。

5.6.2 B+树

考虑如下表结构

```
CREATE TABLE t(
  a INT,
  b INT,
  PRIMARY KEY(a),
  KEY idx_a_b(a,b)
)ENGINE=INNODB
```

B+树索引结构如下

根节点：(2,4)

叶节点1：(1,1) (1,2) (2,1)

叶节点2：(2,4) (3,1) (3,2)

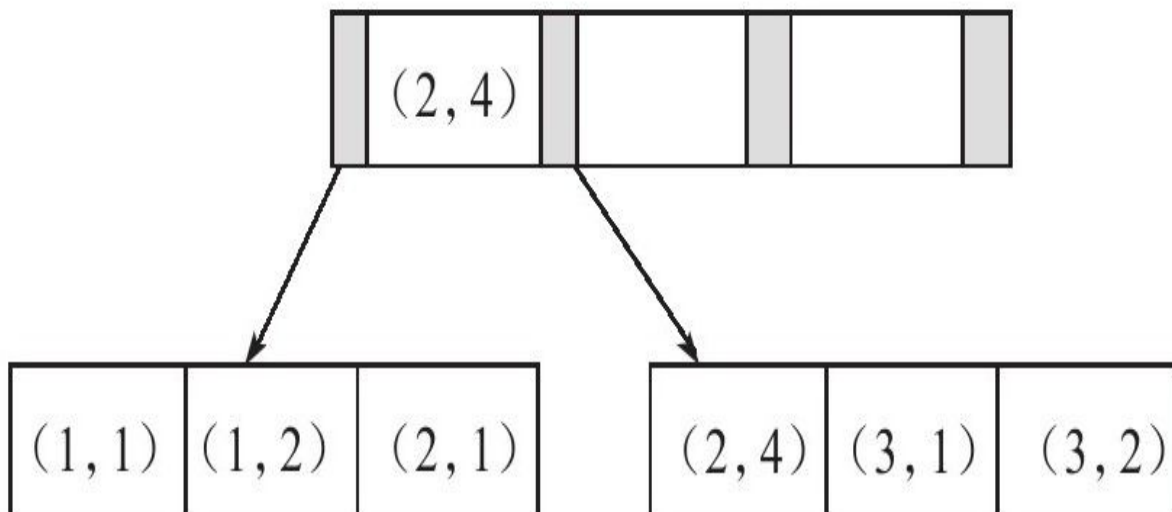


图 5-22 B+树索引结构

5-23

| | id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
|---|----|-------------|---------|------|-----------------|--------|---------|-------|------|-------|
| ▶ | 1 | SIMPLE | buy_log | ref | userid,userid_2 | userid | 4 | const | 1 | |

5-23

possible_keysuserid
useridbuy_dateuserid
userid

userid13SQL5-24

SELECT*FROM buy_log

WHERE userid=1 ORDER BY buy_date DESC LIMIT 3

| | id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
|---|----|-------------|---------|------|-----------------|----------|---------|-------|------|--------------------------|
| ▶ | 1 | SIMPLE | buy_log | ref | userid,userid_2 | userid_2 | 4 | const | 3 | Using where; Using index |

5-24 SQL

SQLuseriduseridbuy_date
useridbuy_dateuserid_2
buy_datebuy_date
userid5-25

| | id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
|---|----|-------------|---------|------|---------------|--------|---------|-------|------|-----------------------------|
| ▶ | 1 | SIMPLE | buy_log | ref | userid | userid | 4 | const | 3 | Using where; Using filesort |

5-25

Extra Using filesort
buy_date userid buy_date

a b a b

```
SELECT...FROM TABLE WHERE a=xxx ORDER BY b
```

a b c

```
SELECT...FROM TABLE WHERE a=xxx ORDER BY b
SELECT...FROM TABLE WHERE a=xxx AND b=xxx ORDER BY c
```

filesort a c

```
SELECT...FROM TABLE WHERE a=xxx ORDER BY c
```

5.6.3 覆盖索引

InnoDB 覆盖索引 (covering index) 是指索引本身包含查询所需的所有数据，因此不需要回表读取数据，从而大大减少了 I/O 次数。

从 InnoDB Plugin 1.0 开始，MySQL 5.0 及以上版本支持 InnoDB 覆盖索引。

InnoDB 覆盖索引的创建语法如下：
primary key1, primary key2, ..., key1, key2, ...

```
SELECT key2 FROM table WHERE key1=xxx
```

```
SELECT primary key2, key2 FROM table WHERE key1=xxx
```

```
SELECT primary key1, key2 FROM table WHERE key1=xxx
```

```
SELECT primary key1, primary key2, key2 FROM table WHERE key1=xxx
```

创建覆盖索引的示例如下：
buy_log

```
SELECT COUNT(*) FROM buy_log;
```

InnoDB 覆盖索引的创建语法如下：
buy_log

| | id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
|---|----|-------------|---------|-------|---------------|--------|---------|------|------|-------------|
| ▶ | 1 | SIMPLE | buy_log | index | NULL | userid | 4 | NULL | 7 | Using index |

图 5-26 COUNT(*) 覆盖索引

5-26 possible_keys NULL userid
Extra Using index

a b b

SELECT COUNT(*)FROM buy_log

WHERE buy_date='2011-01-01'AND buy_date'2011-02-01'

buy_log userid buy_date b
SQL
5-27

| | id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
|--|----|-------------|---------|-------|---------------|----------|---------|------|------|--------------------------|
| | 1 | SIMPLE | buy_log | index | NULL | userid_2 | 8 | NULL | 7 | Using where; Using index |

5-27

5-27 possible_keys NULL key userid_2
userid buy_date Extra Using index

5.6.4 索引的创建

使用EXPLAIN查看SQL语句的执行计划，了解数据库如何执行JOIN操作。

```
SELECT * FROM orderdetails
```

```
WHERE orderid > 10000 and orderid < 102000;
```

使用SHOW INDEX FROM orderdetails查看索引信息。

| | Table | Non_unique | Key_name | Seq_in_index | Column_name | Collation | Cardinality | Sub_part | Packed | Null | Index_type | Comment |
|---|--------------|------------|-----------------------|--------------|-------------|-----------|-------------|----------|--------|------|------------|---------|
| ▶ | orderdetails | 0 | PRIMARY | 1 | OrderID | A | 2311 | NULL | NULL | | BTREE | |
| | orderdetails | 0 | PRIMARY | 2 | ProductID | A | 2311 | NULL | NULL | | BTREE | |
| | orderdetails | 1 | OrderID | 1 | OrderID | A | 2311 | NULL | NULL | | BTREE | |
| | orderdetails | 1 | OrdersOrder_Details | 1 | OrderID | A | 1155 | NULL | NULL | | BTREE | |
| | orderdetails | 1 | ProductID | 1 | ProductID | A | 177 | NULL | NULL | | BTREE | |
| | orderdetails | 1 | ProductsOrder_Details | 1 | ProductID | A | 177 | NULL | NULL | | BTREE | |

图 5-28 orderdetails表的索引

使用EXPLAIN查看SQL语句的执行计划，了解数据库如何执行JOIN操作。

5.6.5 INDEX HINT

MySQL에서 INDEX HINT을 사용하는 방법은 다음과 같다.

MySQL에서 INDEX HINT을 사용하는 방법은 다음과 같다.

SQL에서 INDEX HINT을 사용하는 방법은 다음과 같다.

MySQL에서 Index Hint을 사용하는 방법은 다음과 같다.

```
tbl_name[[AS]alias][index_hint_list]

index_hint_list:

index_hint[,index_hint]...

index_hint:

USE{INDEX|KEY}

[{FOR{JOIN|ORDER BY|GROUP BY}}([index_list])

|IGNORE{INDEX|KEY}

[{FOR{JOIN|ORDER BY|GROUP BY}}(index_list)

|FORCE{INDEX|KEY}

[{FOR{JOIN|ORDER BY|GROUP BY}}(index_list)

index_list:

index_name[,index_name]...
```

다음은 MySQL에서 Index Hint을 사용하는 예제이다.

```
CREATE TABLE t(

a INT,
```

```

b INT,

KEY(a),

KEY(b)

)ENGINE=INNODB;

INSERT INTO t SELECT 1,1;

INSERT INTO t SELECT 1,2;

INSERT INTO t SELECT 2,3;

INSERT INTO t SELECT 2,4;

INSERT INTO t SELECT 1,2;

```

5-30 SQL

```
SELECT*FROM t WHERE a=1 AND b=2;
```

EXPLAIN 5-31

| | id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
|---|----|-------------|-------|-------------|---------------|-----|---------|------|------|--|
| ► | 1 | SIMPLE | t | index_merge | a,b | b,a | 5,5 | NULL | 1 | Using intersect(b,a); Using where; Using index |

5-31 SQL

5-31 possible_keys SQL a b
key a b MySQL a b Extra
Using intersect b a

USE INDEX a

```
SELECT*FROM t USE INDEX(a)WHERE a=1 AND b=2;
```

5-32

5.6.6 Multi-Range Read

MySQL 5.6 Multi-Range Read (MRR) Multi-Range Read IO-bound SQL Multi-Range Read range ref eq_ref

MRR

MRR

InnoDB MyISAM JOIN MRR

RowID

RowID

InnoDB MyISAM SQL

```
SELECT * FROM salaries WHERE salary > 10000 AND salary < 40000;
```

salary idx_s Multi-Range Read 5-34

| | id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
|---|----|-------------|----------|-------|---------------|-------|---------|------|-------|-----------------------|
| ► | 1 | SIMPLE | salaries | range | idx_s | idx_s | 4 | NULL | 23378 | Using index condition |

图 5-34 未启用 Multi-Range Read 的查询结果

图 5-34 未启用 Multi-Range Read 的查询结果。Extra 列显示 Using index condition，表示没有使用 MRR。

| | id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
|---|----|-------------|----------|-------|---------------|-------|---------|------|-------|----------------------------------|
| ► | 1 | SIMPLE | salaries | range | idx_s | idx_s | 4 | NULL | 23378 | Using index condition; Using MRR |

图 5-35 启用 Multi-Range Read 的查询结果

图 5-35 启用 Multi-Range Read 的查询结果。Extra 列显示 Using index condition; Using MRR，表示使用了 MRR。

表 5-4 是否启用 Multi-Range Read 的执行时间对比

| | 执行时间 (秒) |
|----------------------|----------|
| 不使用 Multi-Range Read | 43.213 |
| 使用 Multi-Range Read | 4.212 |

图 5-36 是否启用 Multi-Range Read 的执行时间对比。从图 5-36 可以看出，启用 Multi-Range Read 后，执行时间显著减少。

图 5-36 是否启用 Multi-Range Read 的执行时间对比。从图 5-36 可以看出，启用 Multi-Range Read 后，执行时间显著减少。

Multi-Range Read

```
SELECT*FROM t
WHERE key_part1=1000 AND key_part1<2000
AND key_part2=10000;
```

tkey_part1key_part2key_part1
key_part2Multi-Read RangeRange
SQLkey_part110002000key_part2
1000key_part2Mulit-Range Read

Multi-Range Read
10001000100110001002
1000...19991000

```
SELECT*FROM salaries
WHERE(from_date between'1986-01-01'AND'1995-01-01')
AND(salary between 38000 and 40000);
```

Multi-Range Read5-36

| | id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
|---|----|-------------|----------|-------|---------------|-------|---------|------|--------|----------------------------------|
| ► | 1 | SIMPLE | salaries | range | idx_s | idx_s | 4 | NULL | 210740 | Using index condition; Using MRR |

5-36 Multi-Range Read

salaries salary idx_s SQL Multi-Range Read Extra Using MRR

Multi-Range Read optimizer_switch flag
 mrr on Multi-Range Read mrr_cost_based
 cost based mrr mrr on mrr_cost_based
 off Multi-Range Read Multi-Range
 Read

```
mysql>SET@@optimizer_switch='mrr=on,mrr_cost_based=off';
```

```
Query OK,0 rows affected(0.00 sec)
```

read_rnd_buffer_size
 RowID RowID 256K

```
mysql>SELECT@@read_rnd_buffer_size\G;
```

```
*****1.row*****
```

```
@@read_rnd_buffer_size:262144
```

```
1 row in set(0.00 sec)
```

5.6.7 Index Condition Pushdown (ICP)

Multi-Range Read Index Condition Pushdown MySQL 5.6
MySQL Index Condition Pushdown WHERE
MySQL Index Condition Pushdown MySQL
WHERE WHERE
SQL fetch

Index Condition Pushdown range ref eq_ref ref_or_null
MyISAM InnoDB Index Condition Pushdown
Extra Using index condition

NDB Cluster Engine Condition Pushdown
"Index" Condition Pushdown Condition Pushdown
MySQL 5.1 NDB Cluster Engine Condition Pushdown

(zip_code last_name first_name)

```
SELECT * FROM people  
  
WHERE zipcode='95054'  
  
AND lastname LIKE '%etrunia%'  
  
AND address LIKE '%Main Street%';
```

MySQL zipcode 95 054
WHERE lastname LIKE '%etrunia%' AND address
LIKE '%Main Street%' Index Condition Pushdown
zipcode 95 054 WHERE

Index Condition Pushdown WHERE
WHERE
SQL

```
SELECT*FROM salaries
WHERE(from_date between'1986-01-01'AND'1995-01-01')
AND(salary between 38000 and 40000);
```

Multi-Range Read5-37

| | id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
|--|----|-------------|----------|-------|---------------|-------|---------|------|--------|-----------------------|
| | 1 | SIMPLE | salaries | range | idx_s | idx_s | 4 | NULL | 210740 | Using index condition |

5-37 Multi-Range Read

ExtraUsing index conditionidx_s
Index Condition Pushdown(emp_no
from_date)idx_sfrom_date

5-5MySQL 5.5MySQL 5.6SQL
MRR

表 5-5 MySQL 5.5 和 MySQL 5.6 中是否启用 Index Condition Pushdown 的执行时间对比

| | 执行时间 (秒) |
|--------------------------|------------|
| MySQL 5.5 | 46.738 |
| MySQL 5.6 with ICP | 37.924 |
| MySQL 5.6 with ICP & MRR | 7.816 |

Index Condition Pushdown
MySQL 5.5 23% Mult-Range
Read 400%

5.7 ☐☐☐☐

01
128GB
O1

5.7.1 練習問題

Hash Table U
 $U = \{0, 1, \dots, m-1\}$ ⁽¹⁾

```

0000000000000000T[0..m-1]0000000000000000000000000000U000000
00005-380000000000k000000000000k00000000000000000000k0000
0T[k]=NULL0

```

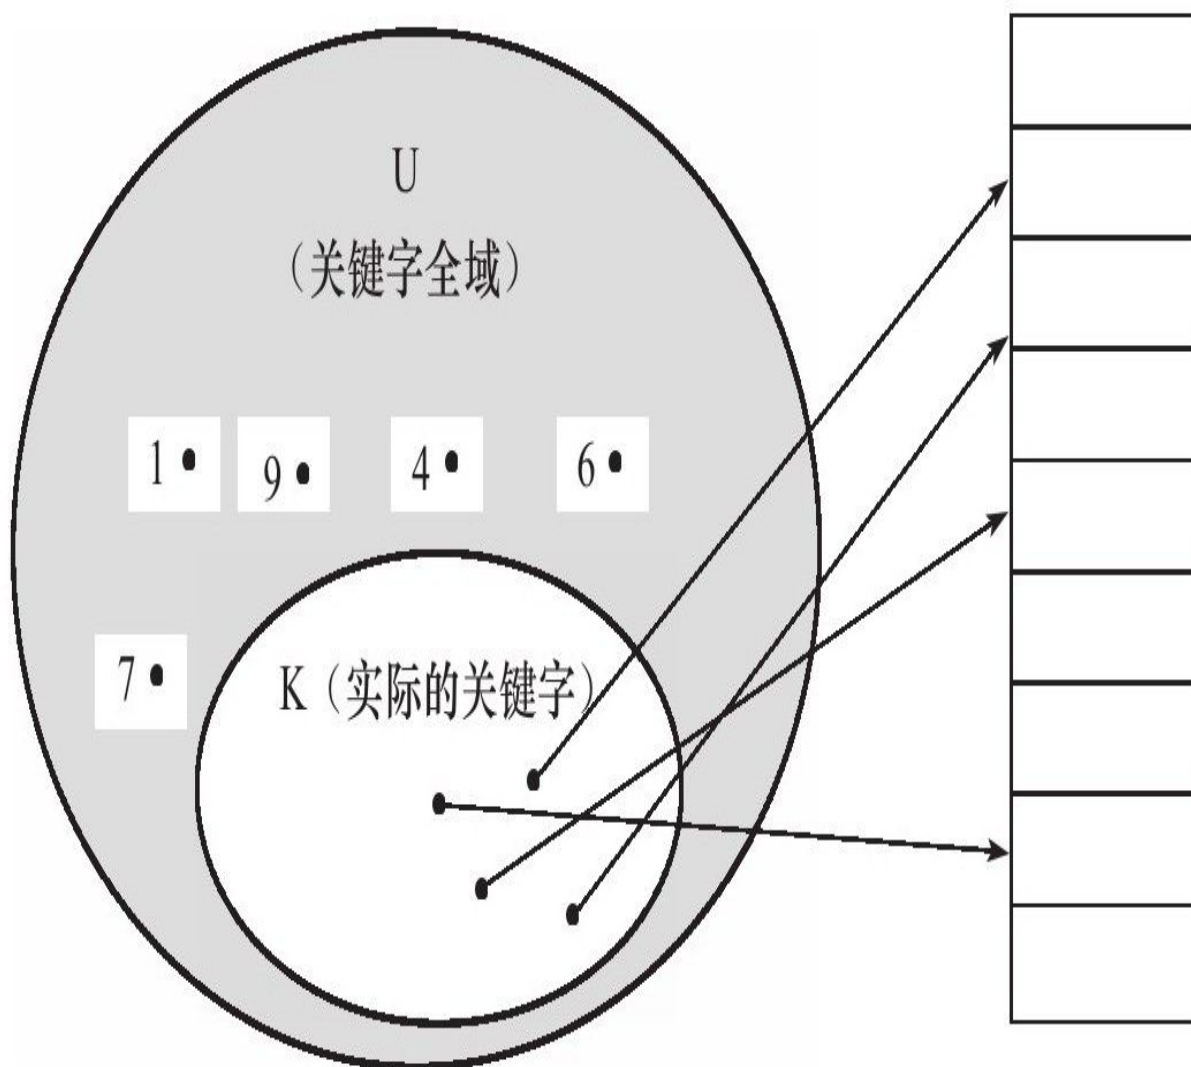


图 5-38 哈希表

哈希表是一种数据结构，它通过哈希函数将关键字映射到数组的索引。哈希表的实现方式有多种，其中一种常见的实现方式是使用数组来存储关键字。在图 5-38 中，哈希表的数组大小为 10，关键字 1、9、4、6 和 7 分别映射到数组的索引 1、2、3、4 和 5。这种实现方式称为“拉链法”。

哈希表的另一种实现方式是使用链表来存储关键字。在图 5-39 中，哈希表的数组大小为 10，关键字 1、9、4、6 和 7 分别映射到数组的索引 1、2、3、4 和 5。这种实现方式称为“桶法”。

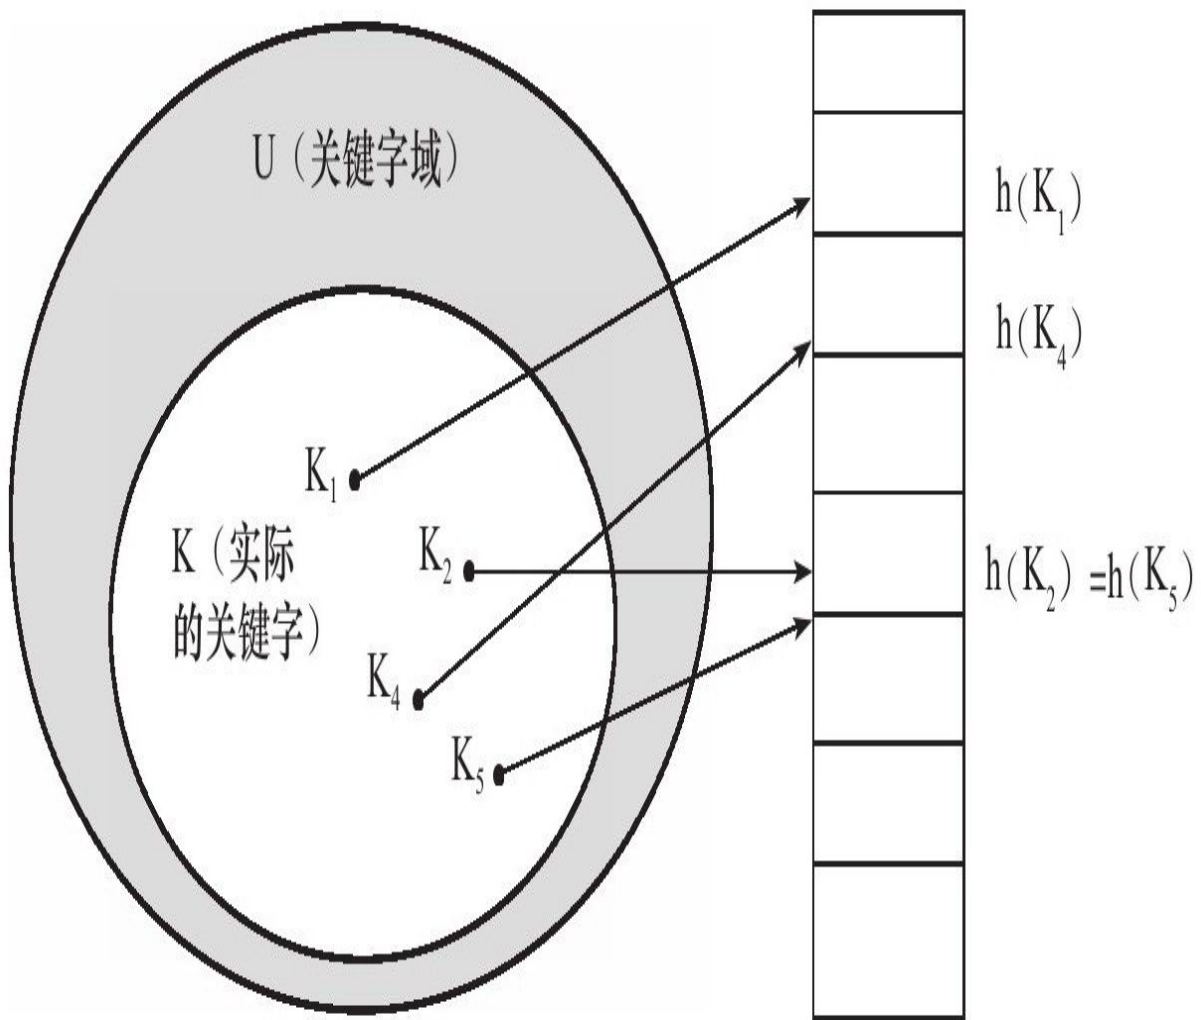


图 5-39 哈希表

哈希表是一种数据结构，它通过哈希函数将关键字映射到表中的槽位。图 5-39 展示了哈希表的实现，其中关键字 K_1, K_2, K_4, K_5 被映射到槽位 $h(K_1), h(K_4), h(K_2), h(K_5)$ 。由于 $h(K_2) = h(K_5)$ ，导致了碰撞（collision）。这种碰撞可以通过链式法（chaining）来解决。

图 5-40 展示了链式法的实现，其中关键字 K_1, K_2, K_4, K_5 被映射到槽位 $h(K_1), h(K_4), h(K_2), h(K_5)$ 。由于 $h(K_2) = h(K_5)$ ，导致了碰撞（collision）。这种碰撞可以通过链式法（chaining）来解决。

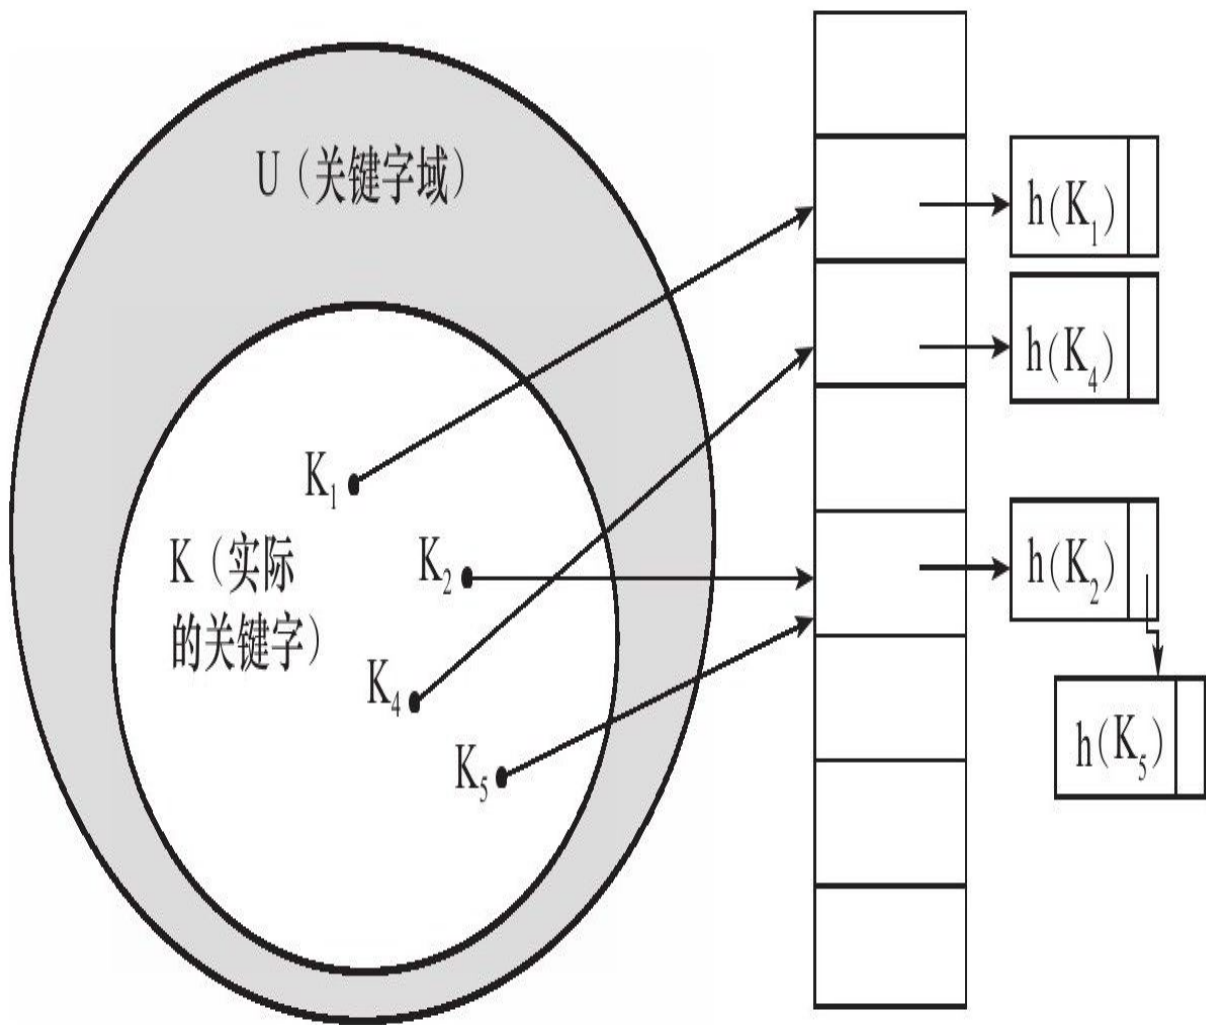


图 5-40 哈希表的示意图

哈希表的构造方法有多种，其中一种常用的方法是直接定址法。直接定址法是指将关键字 k 直接映射到哈希表的地址 $h(k)$ 。这种方法适用于关键字分布比较均匀且没有冲突的情况。

直接定址法的公式为： $h(k) = k$ 。其中 k 是关键字， $h(k)$ 是哈希表的地址。

$$h(k) = k \bmod m$$

[1] m 是哈希表的容量。

5.7.2 InnoDB 数据字典

InnoDB 数据字典存储在 InnoDB 表空间中，由 InnoDB 引擎维护。数据字典表包括 `Page`、`chain`、`m`、`2` 等表。InnoDB 的 `innodb_buffer_pool_size` 默认为 10M，每页 16KB，因此总共有 $640 \times 2 = 1280$ 页。InnoDB 的 `space_id` 默认为 20，因此 `space_id` 为 20 的页的 `offset` 为 1399。

InnoDB 数据字典表存储在 InnoDB 表空间中，由 InnoDB 引擎维护。数据字典表包括 `Page`、`chain`、`m`、`2` 等表。

InnoDB 数据字典表存储在 InnoDB 表空间中，由 InnoDB 引擎维护。数据字典表包括 `space_id`、`16KB`、`offset`、`InnoDB`、`space_id`、`20`、`space_id`、`offset`、`K=space_id`、`20+space_id+offset` 等表。

5.7.3 性能测试

在数据库中执行以下SQL语句，测试数据库性能。DBA用户执行以下SQL语句，测试数据库性能。

```
SELECT*FROM TABLE WHERE index_col='xxx'
```

在数据库中执行以下SQL语句，测试数据库性能。

```
SHOW ENGINE INNODB STATUS
```

```
mysql>SHOW ENGINE INNODB STATUS\G;
```

```
*****1.row*****
```

```
Status:
```

```
=====
```

```
090922 11:52:51 INNODB MONITOR OUTPUT
```

```
=====
```

```
Per second averages calculated from the last 15 seconds
```

```
.....
```

```
-----
```

```
INSERT BUFFER AND ADAPTIVE HASH INDEX
```

```
-----
```

```
Ibuf:size 2249,free list len 3346,seg size 5596,
```

```
374650 inserts,51897 merged recs,14300 merges
```

```
Hash table size 4980499,node heap has 1246 buffer(s)
```

```
1640.60 hash searches/s,3709.46 non-hash searches/s
```

```
.....
```

在数据库中执行以下SQL语句，测试数据库性能。

```
SELECT*FROM table WHERE index_col='xxx'
```

```
SELECT*FROM table WHERE index_col='xxx'
```

non-hash
searches/s hash searches:non-hash searches

InnoDB
innodb_adaptive_hash_index

5.8 B+ 트리

5.8.1 B+ 트리

B+ 트리의 구조는 B 트리와 유사하지만, leaf 노드만 데이터를 저장하며, 모든 leaf 노드가 동일한 크기인 prefix를 공유한다. B+ 트리의 구조는 B 트리와 유사하지만, leaf 노드만 데이터를 저장하며, 모든 leaf 노드가 동일한 크기인 prefix를 공유한다.

```
SELECT * FROM blog WHERE content like 'xxx%'
```

SQL에서 xxx를 content에 B+ 트리를 사용하여 검색한다. B+ 트리의 구조는 B 트리와 유사하지만, leaf 노드만 데이터를 저장하며, 모든 leaf 노드가 동일한 크기인 prefix를 공유한다. xxx를 검색한다.

```
SELECT * FROM blog WHERE content like '%xxx%'
```

B+ 트리의 구조는 B 트리와 유사하지만, leaf 노드만 데이터를 저장하며, 모든 leaf 노드가 동일한 크기인 prefix를 공유한다. B+ 트리의 구조는 B 트리와 유사하지만, leaf 노드만 데이터를 저장하며, 모든 leaf 노드가 동일한 크기인 prefix를 공유한다.

Full-Text Search는 B+ 트리를 사용하여 검색한다. Full-Text Search는 B+ 트리를 사용하여 검색한다. Full-Text Search는 B+ 트리를 사용하여 검색한다.

MySQL은 InnoDB와 MyISAM을 사용한다. InnoDB는 B+ 트리를 사용하며, MyISAM은 B 트리를 사용한다. InnoDB는 B+ 트리를 사용하며, MyISAM은 B 트리를 사용한다.

InnoDB 1.2.x는 InnoDB를 MyISAM으로 대체한다. InnoDB 1.2.x는 InnoDB를 MyISAM으로 대체한다. InnoDB 1.2.x는 InnoDB를 MyISAM으로 대체한다.

5.8.2 倒排索引

倒排索引inverted index和B+树索引auxiliary table类似，都是把索引项按关键字排序，但倒排索引的索引项是倒置的。

倒排文件索引inverted file index{倒排索引项ID}

全倒排索引full inverted index{(倒排索引项ID倒排索引项)}

表5-6 全文检索表t

| DocumentId | Text | DocumentId | Text |
|------------|---|------------|-------------------------------------|
| 1 | Pease porridge hot, pease porridge cold | 4 | Some like it hot, some like it cold |
| 2 | Pease porridge in the pot | 5 | Some like it in the pot |
| 3 | Nine days old | 6 | Nine days old |

DocumentId倒排索引项IDText倒排索引项Some倒排索引项Id倒排索引项Some倒排索引项Id倒排索引项

倒排文件索引inverted file index{倒排索引项ID倒排索引项}5-7

表 5-7 inverted file index 的关联数组

| Number | Text | Documents | Number | Text | Documents |
|--------|------|-----------|--------|----------|-----------|
| 1 | code | 1, 4 | 8 | old | 3, 6 |
| 2 | days | 3, 6 | 9 | pease | 1, 2 |
| 3 | hot | 1, 4 | 10 | porridge | 1, 2 |
| 4 | in | 2, 5 | 11 | pot | 2, 5 |
| 5 | it | 4, 5 | 12 | some | 4, 5 |
| 6 | like | 4, 5 | 13 | the | 2, 5 |
| 7 | nine | 3, 6 | | | |

1 2 3 4 5 6 7 8 9 10 11 12 13
 code days hot in it like nine
 Documents
 inverted file index
 full inverted index (pair) (DocumentId
 Position) 5-8

表 5-8 full inverted index 的关联数组

| Number | Text | Documents | Number | Text | Documents |
|--------|------|----------------|--------|----------|----------------|
| 1 | code | (1:6), (4:8) | 8 | old | (3:3), (6:3) |
| 2 | days | (3:2), (6:2) | 9 | pease | (1:1,4), (2:1) |
| 3 | hot | (1:3), (4:4) | 10 | porridge | (1:2,5), (2:2) |
| 4 | in | (2:3), (5:4) | 11 | pot | (2:5), (5:6) |
| 5 | it | (4:3,7), (5:3) | 12 | some | (4:1,5), (5:1) |
| 6 | like | (4:2,6), (5:2) | 13 | the | (2:4), (5:5) |
| 7 | nine | (3:1), (6:1) | | | |

full inverted index 1:6 1
6 code full inverted index

5.8.3 InnoDB

InnoDB 1.2.x uses a full inverted index. InnoDB (DocumentId, Position) list word list word InnoDB list Position Proximity Search MyISAM

word Auxiliary Table InnoDB Auxiliary Table 6 word Latin

Auxiliary Table InnoDB FTS Index Cache

FTS Index Cache word list FTS Index Cache Auxiliary Table InnoDB Auxiliary Table Auxiliary Table FTS Index Cache word Auxiliary Table merge Insert Buffer Insert Buffer B+ FTS Index Cache Insert Buffer InnoDB Auxiliary Table

InnoDB Auxiliary Table innodb_ft_aux_table Auxiliary Table SQL test_fts_a Auxiliary Table

```
mysql> SET GLOBAL innodb_ft_aux_table='test/fts_a';
```

```
Query OK, 0 rows affected (0.00 sec)
```

```
information_schema
INNODB_FT_INDEX_TABLE
fts_a
```

Oracle 11g에서 InnoDB에서 FTS Index Cache를 InnoDB에서 관리하는 방법을 소개합니다.

FTS Index Cache와 Auxiliary Table은 FTS Index Cache와 Auxiliary Table은 InnoDB에서 관리되는 FTS Index Cache입니다.

innodb_ft_cache_size는 FTS Index Cache의 크기를 32M (word list) Auxiliary Table의 크기를 나타냅니다.

FTS Document ID는 InnoDB에서 FTS_DOC_ID로 저장되며, BIGINT UNSIGNED NOT NULL로 선언됩니다. InnoDB에서 FTS_DOC_ID INDEX Unique Index를 생성하고, FTS_DOC_ID Unique Index를 MySQL에서 생성합니다.

```
mysql> CREATE TABLE fts_a(
- FTS_DOC_ID INT UNSIGNED AUTO_INCREMENT NOT NULL,
- body TEXT,
- PRIMARY KEY(FTS_DOC_ID)
- );
ERROR 1166(42000):Incorrect column name'FTS_DOC_ID'
```

FTS_DOC_ID는 INT로 BIG INT로 선언됩니다. Incorrect column name'FTS_DOC_ID' 오류를 해결합니다.

```
mysql> CREATE TABLE fts_a(
```

```
- FTS_DOC_ID BIGINT UNSIGNED AUTO_INCREMENT NOT NULL,
```

```
- body TEXT,
```

```
- PRIMARY KEY(FTS_DOC_ID)
```

```
- );
```

```
Query OK, 0 rows affected (0.02 sec)
```

MySQL 5.6.23 Auxiliary Table
MySQL 5.6.23 FTS Cache Index Auxiliary Table
MySQL 5.6.23 InnoDB FTS Document ID DELETED
auxiliary table innodb_ft_aux_table
information_schema INNODB_FT_DELETED FTS
Document ID

MySQL 5.6.23 DML DELETED
MySQL 5.6.23 InnoDB
OPTIMIZE TABLE OPTIMIZE TABLE
Cardinality
innodb_optimize_fulltext_only

```
mysql> SET GLOBAL innodb_optimize_fulltext_only=1;
```

```
mysql> OPTIMIZE TABLEfts_a;
```

MySQL 5.6.23 OPTIMIZE TABLE
MySQL 5.6.23
innodb_ft_num_word_optimize
2000

MySQL 5.6.23 fts_a

```
CREATE TABLE fts_a(
```

```
FTS_DOC_ID BIGINT UNSIGNED AUTO_INCREMENT NOT NULL,
```

```
body TEXT,
```

```

PRIMARY KEY(FTS_DOC_ID)

);

INSERT INTO fts_a

SELECT NULL,'Pease porridge in the pot';

INSERT INTO fts_a

SELECT NULL,'Pease porridge hot,pease porridge cold';

INSERT INTO fts_a

SELECT NULL,'Nine days old';

INSERT INTO fts_a

SELECT NULL,'Some like it hot,some like it cold';

INSERT INTO fts_a

SELECT NULL,'Some like it in the pot';

INSERT INTO fts_a

SELECT NULL,'Nine days old';

INSERT INTO fts_a

SELECT NULL,'I like code days';

CREATE FULLTEXT INDEX idx_fts ON fts_a(body);

```

```

CREATE FULLTEXT INDEX idx_fts ON fts_a(body);

```

```
mysql> SELECT * FROM fts_a;
```

```

+-----+-----+
| FTS_DOC_ID | body |
+-----+-----+
| 1 | Pease porridge in the pot |
| 2 | Pease porridge hot,pease porridge cold |
| 3 | Nine days old |
| 4 | Some like it hot,some like it cold |

```

|5|Some like it in the pot|

|6|Nine days old|

|7|I like code days|

+-----+-----+-----+-----+-----+-----+

7 rows in set(0.00 sec)

innodb_ft_aux_table

```
mysql>SET GLOBAL innodb_ft_aux_table='test/fts_a';
```

Query OK,0 rows affected(0.00 sec)

```
mysql>SELECT*FROM information_schema.INNODB_FT_INDEX_TABLE;
```

+-----+-----+-----+-----+-----+-----+

|WORD|FIRST_DOC_ID|LAST_DOC_ID|DOC_COUNT|DOC_ID|POSITION|

+-----+-----+-----+-----+-----+-----+

|code|7|7|1|7|7|

|cold|2|4|2|2|35|

|cold|2|4|2|4|31|

|days|3|7|3|3|5|

|days|3|7|3|6|5|

|days|3|7|3|7|12|

|hot|2|4|2|2|15|

|hot|2|4|2|4|13|

|like|4|7|3|4|5|

|like|4|7|3|4|18|

|like|4|7|3|5|5|

|like|4|7|3|7|2|

|nine|3|6|2|3|0|

|nine|3|6|2|6|0|

|old|3|6|2|3|10|

|old|3|6|2|6|10|

|pease|1|2|2|1|0|

|pease|1|2|2|2|0|

|pease|1|2|2|2|20|

|porridge|1|2|2|1|6|

|porridge|1|2|2|2|6|

|porridge|1|2|2|2|20|

|pot|1|5|2|1|22|

|pot|1|5|2|5|20|

|some|4|5|2|4|0|

|some|4|5|2|4|18|

|some|4|5|2|5|0|

+-----+-----+-----+-----+-----+-----+

27 rows in set(0.00 sec)

word DOC_ID POSITION
FIRST_DOC_ID LAST_DOC_ID DOC_COUNT word
ID ID word

SQL FTS_DOC_ID 7

mysql>DELETE FROM test.fts_a WHERE FTS_DOC_ID=7;

Query OK,1 row affected(0.00 sec)

InnoDB ID
DELETED

mysql>SELECT*FROM INNODB_FT_DELETED;

+-----+

|DOC_ID|

+-----+

|7|

+-----+

1 row in set(0.00 sec)

innodb_ft_deletes ID innodb_ft_deletes
innodb_ft_deletes SQL

mysql>SET GLOBAL innodb_optimize_fulltext_only=1;

Query OK,0 rows affected(0.00 sec)

mysql>OPTIMIZE TABLE test.fts_a;

+-----+-----+-----+-----+

|Table|Op|Msg_type|Msg_text|

+-----+-----+-----+-----+

|test.fts_a|optimize|status|OK|

+-----+-----+-----+-----+

1 row in set(0.01 sec)

mysql>SELECT*FROM INNODB_FT_DELETED;

+-----+

|DOC_ID|

+-----+

|7|

+-----+

1 row in set(0.00 sec)

mysql>SELECT*FROM INNODB_FT_BEING_DELETED;

+-----+

|DOC_ID|

+-----+

|7|

+-----+

1 row in set(0.00 sec)

mysql>OPTIMIZE TABLE test.fts_a;
ERROR 182 (HY000): Invalid InnoDB FTS Doc ID
7
ID
ID

mysql>INSERT INTO test.fts_a SELECT 7, 'I like this days';

ERROR 182 (HY000): Invalid InnoDB FTS Doc ID

stopword
stopword list
word
the
stopword
InnoDB
stopword
information_schema
INNODB_FT_DEFAULT_STOPWORD
36
stopword
innodb_ft_server_stopword_table
stopword

mysql>CREATE TABLE user_stopword(

-value VARCHAR(30)

-)ENGINE=INNODB;

Query OK, 0 rows affected (0.03 sec)

mysql>SET GLOBAL

-innodb_ft_server_stopword_table="test/user_stopword";

Query OK, 0 rows affected (0.00 sec)

InnoDB

delimiter

5.8.4 FULLTEXT

MySQL FULLTEXT Search

```
MATCH(col1,col2,...)AGAINST(expr[search_modifier])
```

```
search_modifier:
```

```
{
```

```
IN NATURAL LANGUAGE MODE
```

```
|IN NATURAL LANGUAGE MODE WITH QUERY EXPANSION
```

```
|IN BOOLEAN MODE
```

```
|WITH QUERY EXPANSION
```

```
}
```

MySQL MATCH()...AGAINST() FULLTEXT Search MATCH AGAINST

1.Natural Language

MySQL MATCH()...AGAINST() FULLTEXT Search Natural Language word 5.8.3 fts_a body Pease SQL

```
mysql>SELECT*FROM fts_a WHERE body LIKE'%Pease%';
```

MySQL B+ SQL

```
mysql>SELECT*FROM fts_a
```

```
->WHERE MATCH(body)
```

```
->AGAINST('Porridge'IN NATURAL LANGUAGE MODE);
```

```
+-----+-----+-----+-----+-----+-----+
```

-----+

|1|Pease porridge in the pot|

+

| | |
|--|---|
| | <div style="border-bottom: 1px solid black; width: 60%;"></div> |
|--|---|

[illegible]

[illegible]

```
- WHERE MATCH(body) AGAINST ('Porridge') \G;
```

```
*****1.row*****
```

```
select_type:SIMPLE
```

type:fulltext

```
key:idx_fts
```

ref:NULL

Extra:Using where

[illegible]

type fulltext key
idx_fts
MATCH

```
mysql> SELECT * FROM fts_b
```

```
-> WHERE MATCH(body) AGAINST('Porridge');
```

```
ERROR 1191 (HY000): Can't find FULLTEXT index matching the column list
```

WHERE MATCH Relevance
0
MySQL

word

word

word

word

Porridge 2

MATCH SQL

```
mysql> SELECT count(*)
```

```
-> FROM fts_a WHERE
```

```
-> MATCH(body) AGAINST('Porridge' IN NATURAL LANGUAGE MODE);
```

```
+-----+
```

```
|count(FTS_DOC_ID)|
```

```
+-----+
```

```
|2|
```

```
+-----+
```

```
1 row in set (0.00 sec)
```

SQL

```
mysql> SELECT
- COUNT(IF(MATCH(body)
- AGAINST('Porridge' IN NATURAL LANGUAGE MODE),1,NULL))
- AS count
- FROM fts_a;

+-----+
|count|
+-----+
|2|
+-----+

1 row in set (0.00 sec)
```

SQL

SQL

SQL

SQL

```
mysql> SELECT fts_doc_id,body,
- MATCH(body)AGAINST('Porridge' IN NATURAL LANGUAGE MODE)
- AS Relevance
- FROM fts_a;

+-----+-----+-----+
|fts_doc_id|body|Relevance|
+-----+-----+-----+
|1|Pease porridge in the pot|0.2960100471973419|
|2|Pease porridge hot,pease porridge cold|0.5920200943946838|
|3|Nine days old|0|
```

```
|4|Some like it hot,some like it cold|0|
```

```
|5|Some like it in the pot|0|
```

```
|6|Nine days old|0|
```

```
|7|I like hot and code days|0|
```

```
+-----+-----+-----+-----+
```

```
7 rows in set(0.01 sec)
```

MySQL InnoDB 全文索引

创建 word 和 stopwords 表

```
CREATE TABLE word (
  id INT(11) UNSIGNED NOT NULL AUTO_INCREMENT,
  text VARCHAR(255) NOT NULL,
  innodb_ft_min_token_size INT(11) UNSIGNED NOT NULL,
  innodb_ft_max_token_size INT(11) UNSIGNED NOT NULL,
  PRIMARY KEY (id)
) ENGINE=InnoDB;
```

创建 stopwords 表

```
mysql> SELECT fts_doc_id AS id,body,
  - MATCH(body)AGAINST('the' IN NATURAL LANGUAGE MODE)
  - AS rl
  - FROM fts_a;
```

```
+-----+-----+-----+-----+
```

```
|id|body|rl|
```

```
+-----+-----+-----+-----+
```

```
|1|Pease porridge in the pot|0|
```

```
|2|Pease porridge hot,pease porridge cold|0|
```

```
|3|Nine days old|0|
```

```
|4|Some like it hot,some like it cold|0|
```

```
|5|Some like it in the pot|0|
```

```
|6|Nine days old|0|
```

```
|7|I like hot and code days|0|
```

```
+-----+-----+-----+-----+
```

7 rows in set(0.00 sec)

the 15 stopwords 0

innodb_ft_min_token_size innodb_ft_max_token_size
InnoDB innodb_ft_min_token_size
innodb_ft_max_token_size InnoDB
innodb_ft_min_token_size 3
innodb_ft_max_token_size 84

2.Boolean

MySQL IN BOOLEAN MODE
Pease hot +
-

```
mysql> SELECT * FROM fts_a
```

```
- WHERE MATCH(body) AGAINST ('+Pease-hot' IN BOOLEAN MODE) \G;
```

```
*****1. row*****
```

```
FTS_DOC_ID:1
```

```
body:Pease porridge in the pot
```

Boolean

+ word

- word

no operator word

@distance distance distance
Proximity Search MATCH body AGAINST
"Pease pot"@30 IN BOOLEAN MODE Pease pot
30


```
- WHERE MATCH(body)AGAINST('Pease hot'IN BOOLEAN MODE);
```

```
+-----+-----+
```

```
|FTS_DOC_ID|body|
```

```
+-----+-----+
```

```
|2|Pease porridge hot,pease porridge cold|
```

```
|1|Pease porridge in the pot|
```

```
|4|Some like it hot,some like it cold|
```

```
|7|I like hot and code days|
```

```
+-----+-----+
```

```
4 rows in set(0.00 sec)
```

SQL Proximity Search

```
mysql>SELECT fts_doc_id,body FROM fts_a
```

```
- WHERE MATCH(body)
```

```
- AGAINST('"Pease pot"@30'IN BOOLEAN MODE)\G;
```

```
*****1.row*****
```

```
fts_doc_id:1
```

```
body:Pease porridge in the pot
```

```
1 row in set(0.01 sec)
```

```
mysql>SELECT fts_doc_id,body FROM fts_a
```

```
- WHERE MATCH(body)
```

```
- AGAINST('"Pease pot"@10'IN BOOLEAN MODE);
```

```
Empty set(0.01 sec)
```

1Peasepot22@30
@10

```
mysql>SELECT fts_doc_id,body,
```

```
- MATCH(body)AGAINST('likepot'IN BOOLEAN MODE)
```



```
- AS Relevance FROM fts_a;
```

```
+-----+-----+-----+
```

```
|fts_doc_id|body|Relevance|
```

```
+-----+-----+-----+
```

```
|1|Pease porridge in the pot|1.2960100173950195|
```

```
|2|Pease porridge hot,pease porridge cold|0|
```

```
|3|Nine days old|0|
```

```
|4|Some like it hot,some like it cold|0.27081382274627686|
```

```
|5|Some like it in the pot|1.4314169883728027|
```

```
|6|Nine days old|0|
```

```
|7|I like hot and code days|0.13540691137313843|
```

```
+-----+-----+-----+
```

```
7 rows in set(0.00 sec)
```

SQL like pot pot
4 like pot 1 5

“some”

```
mysql> SELECT fts_doc_id,body,
```

```
- MATCH(body)AGAINST('like hot some'IN BOOLEAN MODE)
```

```
- AS Relevance
```

```
- FROM fts_a;
```

```
+-----+-----+-----+
```

```
|fts_doc_id|body|Relevance|
```

```
+-----+-----+-----+
```

```
|1|Pease porridge in the pot|0|
```

```
|2|Pease porridge hot,pease porridge cold|1.2960100173950195|
```

```
|3|Nine days old|0|
```

```
|4|Some like it hot,some like it cold|1.158843994140625|
```

```
|5|Some like it in the pot|-0.5685830116271973|
```

```
|6|Nine days old|0|
```

```
|7|I like hot and code days|0.13540691137313843|
```

```
+-----+-----+-----+-----+
```

```
7 rows in set(0.00 sec)
```

5 like some

SQL

```
mysql>SELECT*FROM fts_a
```

```
->WHERE MATCH(body)AGAINST('po*'IN BOOLEAN MODE);
```

```
+-----+-----+-----+-----+
```

```
|FTS_DOC_ID|body|
```

```
+-----+-----+-----+-----+
```

```
|2|Pease porridge hot,pease porridge cold|
```

```
|1|Pease porridge in the pot|
```

```
|5|Some like it in the pot|
```

```
+-----+-----+-----+-----+
```

```
3 rows in set(0.00 sec)
```

po porridge pot

SQL

```
mysql>SELECT*FROM fts_a
```

```
->WHERE MATCH(body)AGAINST('like hot'IN BOOLEAN MODE);
```

```
+-----+-----+-----+-----+
```

```
|FTS_DOC_ID|body|
```

```
+-----+-----+-----+-----+
```

```

|4|Some like it hot,some like it cold|

|7|I like hot and code days|

|2|Pease porridge hot,pease porridge cold|

|5|Some like it in the pot|

+-----+-----+

4 rows in set(0.00 sec)

mysql>SELECT*FROM fts_a

- WHERE MATCH(body)AGAINST('"like hot"'IN BOOLEAN MODE);

+-----+-----+

|FTS_DOC_ID|body|

+-----+-----+

|7|I like hot and code days|

+-----+-----+

1 row in set(0.00 sec)

```

MySQL ""like hot "" 4 "" SQL ""like hot"" 4

3.Query Expansion

MySQL implied knowledge database MySQL Oracle DB2 RDBMS Query Expansion implied knowledge

WITH QUERY EXPANSION IN NATURAL LANGUAGE MODE WITH QUERY EXPANSION blind query expansion automatic relevance feedback

articles

```

CREATE TABLE articles(

id INT UNSIGNED AUTO_INCREMENT NOT NULL PRIMARY KEY,

title VARCHAR(200),

body TEXT,

FULLTEXT(title,body)

)ENGINE=InnoDB;

INSERT INTO articles(title,body)VALUES

('MySQL Tutorial','DBMS stands for DataBase...'),

('How To Use MySQL Well','After you went through a...'),

('Optimizing MySQL','In this tutorial we will show...'),

('1001 MySQL Tricks','1.Never run mysqld as root.2....'),

('MySQL vs.YourSQL','In the following database comparison...'),

('MySQL Security','When configured properly,MySQL...'),

('Tuning DB2','For IBM database...'),

('IBM History','DB2 hitory for IBM...');

```

```

mysql>CREATE FULLTEXT INDEX FTS_DOC_ID ON InnoDB database
articles(title,body)

```

```

mysql>SELECT*FROM articles

- WHERE MATCH(title,body)

- AGAINST('database' IN NATURAL LANGUAGE MODE);

+---+-----+-----+
|id|title|body|
+---+-----+-----+
|1|MySQL Tutorial|DBMS stands for DataBase...|
|5|MySQL vs.YourSQL|In the following database comparison...|
|7|Tuning DB2|For IBM database...|

```


5.9 B+

B+ InnoDB 1.2 B+
InnoDB 1.2 B+
InnoDB 1.2 B+

第6章

本章主要介绍数据库的锁机制。首先介绍数据库锁的概念，然后介绍数据库锁的分类，接着介绍数据库锁的实现原理，最后介绍数据库锁的应用。本章的重点是数据库锁的实现原理，包括锁的分配、持有、释放和升级等。本章的难点是数据库锁的实现原理，包括锁的分配、持有、释放和升级等。

本章主要介绍数据库的锁机制。首先介绍数据库锁的概念，然后介绍数据库锁的分类，接着介绍数据库锁的实现原理，最后介绍数据库锁的应用。本章的重点是数据库锁的实现原理，包括锁的分配、持有、释放和升级等。本章的难点是数据库锁的实现原理，包括锁的分配、持有、释放和升级等。

6.1 □□□□

```

InnoDB
InnoDB
LRU
LRU

```

Database Management Systems

- SQL
- MySQL
- InnoDB
- MyISAM
- NDB Cluster
- Microsoft SQL Server
- Oracle

```

MyISAM
“”MyISAM
Microsoft SQL
Server
Microsoft SQL Server 2005
MyISAM
2005
Microsoft SQL Server
InnoDB
Microsoft SQL
Server

```

InnoDB Oracle
 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111 112 113 114 115 116 117 118 119 120 121 122 123 124 125 126 127 128 129 130 131 132 133 134 135 136 137 138 139 140 141 142 143 144 145 146 147 148 149 150 151 152 153 154 155 156 157 158 159 160 161 162 163 164 165 166 167 168 169 170 171 172 173 174 175 176 177 178 179 180 181 182 183 184 185 186 187 188 189 190 191 192 193 194 195 196 197 198 199 200 201 202 203 204 205 206 207 208 209 210 211 212 213 214 215 216 217 218 219 220 221 222 223 224 225 226 227 228 229 230 231 232 233 234 235 236 237 238 239 240 241 242 243 244 245 246 247 248 249 250 251 252 253 254 255 256 257 258 259 260 261 262 263 264 265 266 267 268 269 270 271 272 273 274 275 276 277 278 279 280 281 282 283 284 285 286 287 288 289 290 291 292 293 294 295 296 297 298 299 300 301 302 303 304 305 306 307 308 309 310 311 312 313 314 315 316 317 318 319 320 321 322 323 324 325 326 327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 378 379 380 381 382 383 384 385 386 387 388 389 390 391 392 393 394 395 396 397 398 399 400 401 402 403 404 405 406 407 408 409 410 411 412 413 414 415 416 417 418 419 420 421 422 423 424 425 426 427 428 429 430 431 432 433 434 435 436 437 438 439 440 441 442 443 444 445 446 447 448 449 450 451 452 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476 477 478 479 480 481 482 483 484 485 486 487 488 489 490 491 492 493 494 495 496 497 498 499 500 501 502 503 504 505 506 507 508 509 510 511 512 513 514 515 516 517 518 519 520 521 522 523 524 525 526 527 528 529 530 531 532 533 534 535 536 537 538 539 540 541 542 543 544 545 546 547 548 549 550 551 552 553 554 555 556 557 558 559 560 561 562 563 564 565 566 567 568 569 570 571 572 573 574 575 576 577 578 579 580 581 582 583 584 585 586 587 588 589 590 591 592 593 594 595 596 597 598 599 600 601 602 603 604 605 606 607 608 609 610 611 612 613 614 615 616 617 618 619 620 621 622 623 624 625 626 627 628 629 630 631 632 633 634 635 636 637 638 639 640 641 642 643 644 645 646 647 648 649 650 651 652 653 654 655 656 657 658 659 660 661 662 663 664 665 666 667 668 669 670 671 672 673 674 675 676 677 678 679 680 681 682 683 684 685 686 687 688 689 690 691 692 693 694 695 696 697 698 699 700 701 702 703 704 705 706 707 708 709 710 711 712 713 714 715 716 717 718 719 720 721 722 723 724 725 726 727 728 729 730 731 732 733 734 735 736 737 738 739 740 741 742 743 744 745 746 747 748 749 750 751 752 753 754 755 756 757 758 759 760 761 762 763 764 765 766 767 768 769 770 771 772 773 774 775 776 777 778 779 780 781 782 783 784 785 786 787 788 789 790 791 792 793 794 795 796 797 798 799 800 801 802 803 804 805 806 807 808 809 810 811 812 813 814 815 816 817 818 819 820 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836 837 838 839 840 841 842 843 844 845 846 847 848 849 850 851 852 853 854 855 856 857 858 859 860 861 862 863 864 865 866 867 868 869 870 871 872 873 874 875 876 877 878 879 880 881 882 883 884 885 886 887 888 889 890 891 892 893 894 895 896 897 898 899 900 901 902 903 904 905 906 907 908 909 910 911 912 913 914 915 916 917 918 919 920 921 922 923 924 925 926 927 928 929 930 931 932 933 934 935 936 937 938 939 940 941 942 943 944 945 946 947 948 949 950 951 952 953 954 955 956 957 958 959 960 961 962 963 964 965 966 967 968 969 970 971 972 973 974 975 976 977 978 979 980 981 982 983 984 985 986 987 988 989 990 991 992 993 994 995 996 997 998 999 1000 1001 1002 1003 1004 1005 1006 1007 1008 1009 1010 1011 1012 1013 1014 1015 1016 1017 1018 1019 1020 1021 1022 1023 1024 1025 1026 1027 1028 1029 1030 1031 1032 1033 1034 1035 1036 1037 1038 1039

[1] _____ “_____” _____ “_____” _____

6.2 lock_latch

```
lock latch lock latch " "
lock
```

```
latch
InnoDB latch mutex rwlock
```

```
lock[00000000000000000000000000000000]lock[000000]
commit[rollback[00000000000000000000000000000000]lock[000000]
0000000000000000006-1[000]lock[latch[000]
```

表 6-1 lock 与 latch 的比较

| | lock | latch |
|------|--|--|
| 对象 | 事务 | 线程 |
| 保护 | 数据库内容 | 内存数据结构 |
| 持续时间 | 整个事务过程 | 临界资源 |
| 模式 | 行锁、表锁、意向锁 | 读写锁、互斥量 |
| 死锁 | 通过 waits-for graph、time out 等机制进行死锁检测与处理 | 无死锁检测与处理机制。仅通过应用程序加锁的顺序 (lock leveling) 保证无死锁的情况发生 |
| 存在于 | Lock Manager 的哈希表中 | 每个数据结构的对象中 |

❑❑InnoDB❑❑❑❑❑latch❑❑❑❑❑❑SHOW ENGINE INNODB MUTEX
❑❑❑❑❑❑❑6-1❑❑❑

```
mysql> SHOW ENGINE INNODB MUTEX;
```

| Type | Name | Status |
|--------|----------------|------------|
| InnoDB | srv0srv.c:1020 | os_waits=5 |
| InnoDB | log0log.c:833 | os_waits=3 |

2 rows in set (0.03 sec)

❑ 6-1 ❑❑❑❑SHOW ENGINE INNODB MUTEX❑❑latch

❑Debug❑❑❑❑❑❑❑SHOW ENGINE INNODB MUTEX❑❑❑❑latch❑❑❑
❑❑❑❑❑6-2❑❑❑

```
mysql> SHOW ENGINE INNODB MUTEX;
```

| Type | Name | Status |
|--------|-------------------------|--|
| InnoDB | &kernel_mutex:srv0srv.c | count=54, spin_waits=6, spin_rounds=60, os_waits=3, os_yields=3, os_wait_times=0 |
| InnoDB | log0log.c:833 | os_waits=2 |
| InnoDB | rw_lock_mutexes | count=0, spin_waits=0, spin_rounds=0, os_waits=0, os_yields=0, os_wait_times=0 |

3 rows in set (0.01 sec)

latchlockSHOW ENGINE
INNODB STATUSinformation_schemaINNODB_TRX
INNODB_LOCKSINNODB_LOCK_WAITS

6.3 InnoDB锁兼容性

6.3.1 锁类型

InnoDB锁类型分为以下几种：

□共享锁 S Lock

□排他锁 X Lock

假设事务T1对数据r加锁，事务T2对数据r加锁，如果T2加的锁与T1加的锁兼容，则T2可以加锁，否则T2必须等待T1释放锁。Lock Compatible表示两个锁是否兼容。T1对T2加的锁——兼容性表6-3所示。

表 6-3 排他锁和共享锁的兼容性

| | X | S |
|---|-----|-----|
| X | 不兼容 | 不兼容 |
| S | 不兼容 | 兼容 |

表6-3中X对X不兼容，S对S兼容，S对X兼容。row表示行级锁。

InnoDB锁粒度为granular，即行级锁。InnoDB锁粒度为fine granularity，即细粒度锁。表6-3所示。

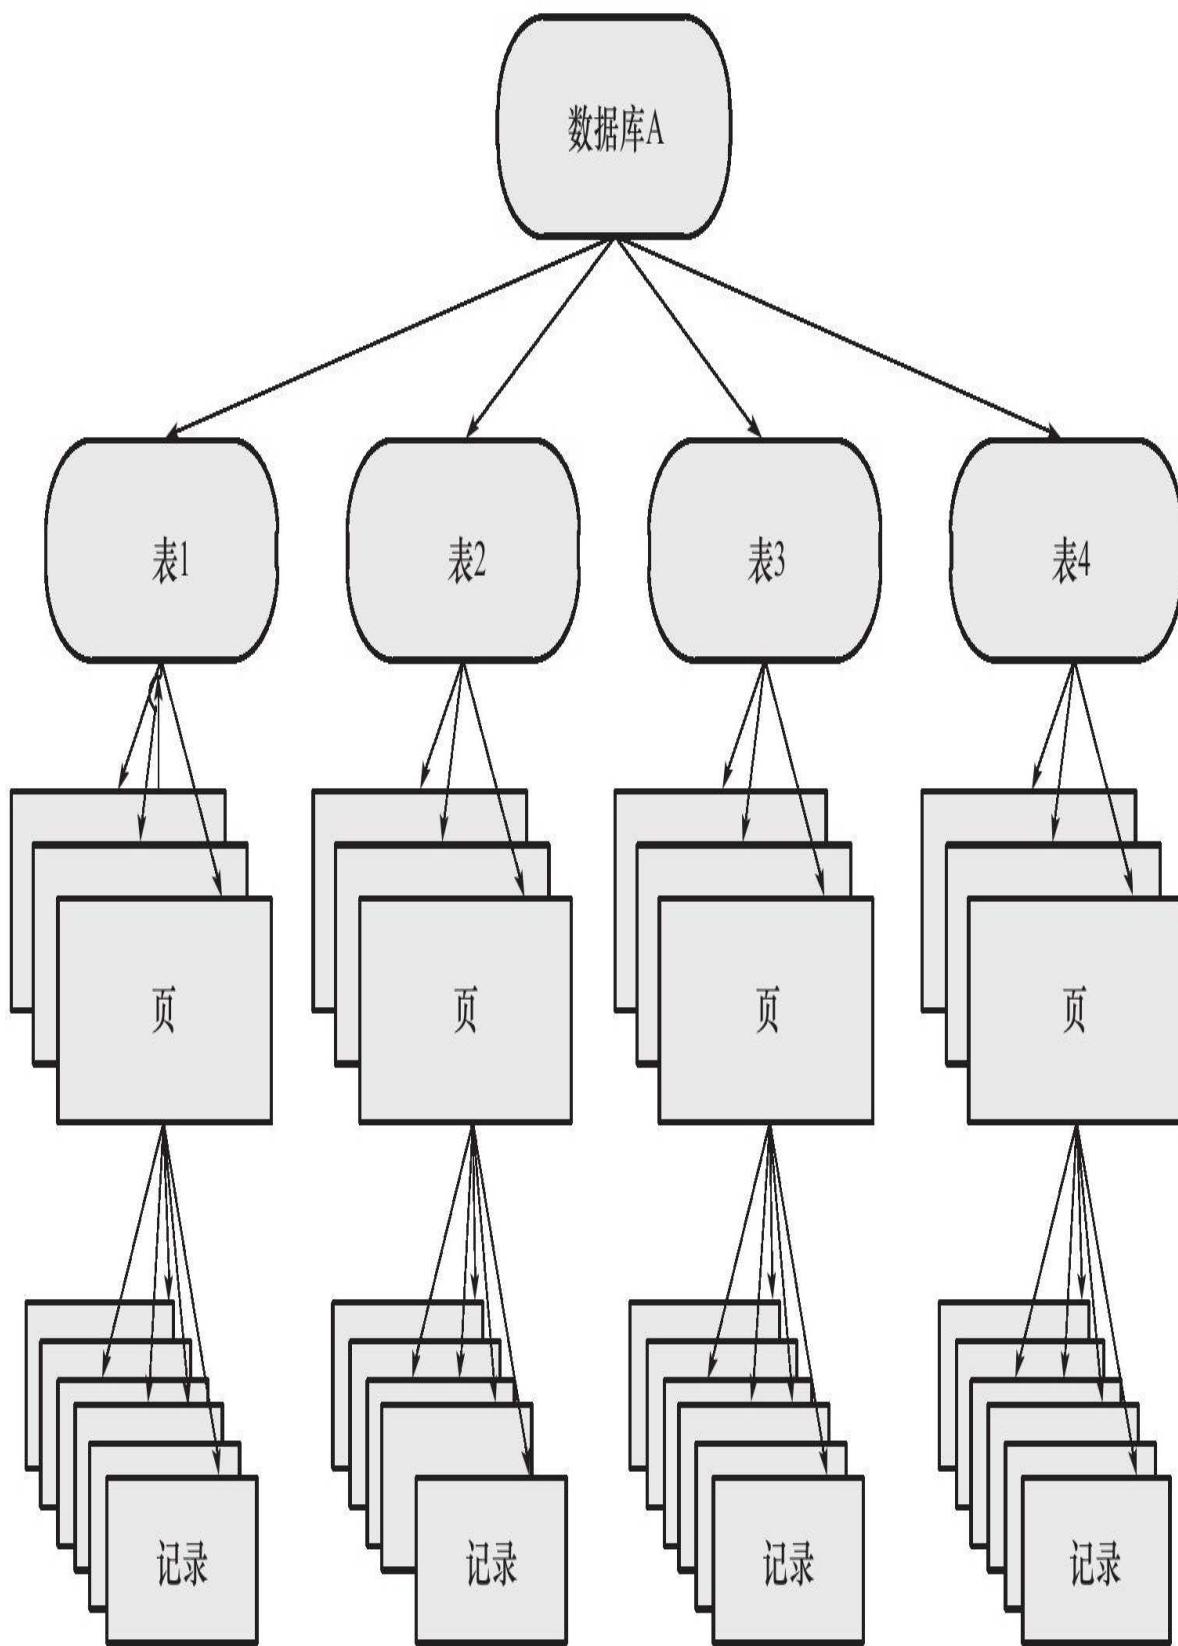


图 6-3 兼容性

图 6-3 展示了 InnoDB 存储引擎中锁的兼容性。图中列出了 IS、IX、S 和 X 四种锁类型，并显示了它们之间的兼容性。图中还列出了 InnoDB 存储引擎中锁的兼容性。图中列出了 InnoDB 存储引擎中锁的兼容性。

InnoDB 存储引擎中锁的兼容性。图中列出了 InnoDB 存储引擎中锁的兼容性。

1. IS Lock 兼容性

2. IX Lock 兼容性

InnoDB 存储引擎中锁的兼容性。图中列出了 InnoDB 存储引擎中锁的兼容性。

表 6-4 InnoDB 存储引擎中锁的兼容性

| | IS | IX | S | X |
|----|-----|-----|-----|-----|
| IS | 兼容 | 兼容 | 兼容 | 不兼容 |
| IX | 兼容 | 兼容 | 不兼容 | 不兼容 |
| S | 兼容 | 不兼容 | 兼容 | 不兼容 |
| X | 不兼容 | 不兼容 | 不兼容 | 不兼容 |

通过执行 SHOW ENGINE INNODB STATUS 命令可以查看 InnoDB 存储引擎中锁的兼容性。

```
mysql>SHOW ENGINE INNODB STATUS\G;
```

```
.....
```

```
-----
```

```
TRANSACTIONS
```

```
-----
```

```
Trx id counter 48B89BF
```

```
Purge done for trx's n:o 48B89BA undo n:o 0
```

```
History list length 0
```

```
LIST OF TRANSACTIONS FOR EACH SESSION:
```

```
---TRANSACTION 0,not started,process no 13757,OS thread id 1255176512
```

```
MySQL thread id 42,query id 80424887 localhost root
```

```
show engine innodb status
```

```
---TRANSACTION 48B89BE,ACTIVE 193 sec,process no 13757,OS thread id 1254910272 starting index read
```

```
mysql tables in use 1,locked 1
```

```
LOCK WAIT 2 lock struct(s),heap size 368,1 row lock(s)
```

```
MySQL thread id 41,query id 80424886 localhost root Sending data
```

```
select*from t where a=4 lock in share mode
```

```
-----TRX HAS BEEN WAITING 2 SEC FOR THIS LOCK TO BE GRANTED:
```

```
RECORD LOCKS space id 30 page no 3 n bits 72 index'PRIMARY'of table'test'.'t'trx id 48B89BE lock mode S waiting
```

```
-----
```

```
TABLE LOCK table'test'.'t'trx id 48B89BE lock mode IS
```

```
RECORD LOCKS space id 30 page no 3 n bits 72 index'PRIMARY'of table'test'.'t'trx id 48B89BE lock mode S waiting
```

```
---TRANSACTION 48B89BD,ACTIVE 205 sec,process no 13757,OS thread id 1257838912
```

```
2 lock struct(s),heap size 368,1 row lock(s)
```

```
MySQL thread id 40,query id 80424881 localhost root
```

```
TABLE LOCK table'test'.'t'trx id 48B89BD lock mode IX
```

```
RECORD LOCKS space id 30 page no 3 n bits 72 index'PRIMARY'of table'test'.'t'trx id 48B89BD lock_mode X locks rec but not gap
```

```
-----
```

```
END OF INNODB MONITOR OUTPUT
```


=====

1 row in set(0.01 sec)

SQL select*from t where a=4 lock in share mode
RECORD LOCKS space id 30 page no 3 n bits 72
index'PRIMARY'of table'test'.'t'trx id 48B89BD lock_mode X
locks rec but not gaplocks rec but not gap
locks rec but not gap

InnoDB 1.0 SHOW FULL PROCESSLIST SHOW
ENGINE INNODB STATUS
InnoDB1.0 INFORMATION_SCHEMA INNODB_TRX
INNODB_LOCKS INNODB_LOCK_WAITS
INNODB_TRX8

表 6-5 表 INNODB_TRX 的结构说明

| 字段名 | 说明 |
|-----------------------|--|
| trx_id | InnoDB 存储引擎内部唯一的事务 ID |
| trx_state | 当前事务的状态 |
| trx_started | 事务的开始时间 |
| trx_requested_lock_id | 等待事务的锁 ID。如 trx_state 的状态为 LOCK WAIT，那么该值代表当前的事务等待之前事务占用锁资源的 ID。若 trx_state 不是 LOCK WAIT，则该值为 NULL |
| trx_wait_started | 事务等待开始的时间 |
| trx_weight | 事务的权重，反映了一个事务修改和锁住的行数。在 InnoDB 存储引擎中，当发生死锁需要回滚时，InnoDB 存储引擎会选择该值最小的进行回滚 |
| trx_mysql_thread_id | MySQL 中的线程 ID，SHOW PROCESSLIST 显示的结果 |
| trx_query | 事务运行的 SQL 语句 |

□□□□□□□□□□

```
mysql>SELECT*FROM information_schema.INNODB_TRX\G;

*****1.row*****

trx_id:7311F4
```

```
trx_state:LOCK WAIT

trx_started:2010-01-04 10:49:33

trx_requested_lock_id:7311F4:96:3:2

trx_wait_started:2010-01-04 10:49:33

trx_weight:2

trx_mysql_thread_id:471719

trx_query:select*from parent lock in share mode

*****2.row*****

trx_id:730FEE

trx_state:RUNNING

trx_started:2010-01-04 10:18:37

trx_requested_lock_id:NULL

trx_wait_started:NULL

trx_weight:2

trx_mysql_thread_id:471718

trx_query:NULL

2 rows in set(0.00 sec)
```

state trx_id 730FEE trx_id 7311F4
“LOCK WAIT” SQL select*from parent lock in
share mode InnoDB
INNODB_LOCKS 6-6

表 6-6 表 INNODB_LOCKS 的结构

| 字段名 | 说明 |
|-------------|---------------------------|
| lock_id | 锁的 ID |
| lock_trx_id | 事务 ID |
| lock_mode | 锁的模式 |
| lock_type | 锁的类型，表锁还是行锁 |
| lock_table | 要加锁的表 |
| lock_index | 锁住的索引 |
| lock_space | 锁对象的 space id |
| lock_page | 事务锁定页的数量。若是表锁，则该值为 NULL |
| lock_rec | 事务锁定行的数量，若是表锁，则该值为 NULL |
| lock_data | 事务锁定记录的主键值，若是表锁，则该值为 NULL |

查看表INNODB_LOCKS

```
mysql>SELECT*FROM information_schema.INNODB_LOCKS\G;

*****1.row*****

lock_id:7311F4:96:3:2

lock_trx_id:7311F4
```

```

lock_mode:S

lock_type:RECORD

lock_table:'mytest'.'parent'

lock_index:'PRIMARY'

lock_space:96

lock_page:3

lock_rec:2

lock_data:1

*****2.row*****

lock_id:730FEE:96:3:2

lock_trx_id:730FEE

lock_mode:X

lock_type:RECORD

lock_table:'mytest'.'parent'

lock_index:'PRIMARY'

lock_space:96

lock_page:3

lock_rec:2

lock_data:1

2 rows in set(0.00 sec)

```

트랜잭션 ID 730FEE가 parent 트랜잭션의 X锁을 ID 7311F4가 parent 트랜잭션의 S锁을 lock_data 1에 락을 걸었다. InnoDB_TRX 테이블에서 trx_state가 "RUNNING"이고 "LOCK WAIT" 상태이다.

lock_data 1에 락을 걸었다. InnoDB_TRX 테이블에서 trx_state가 "RUNNING"이고 "LOCK WAIT" 상태이다. InnoDB_LOCKS 테이블에서 trx_id가 NULL인 InnoDB 트랜잭션이 있다.

表 INNO_DB_LOCKS 是 MySQL 数据库系统内部使用的一个表，它记录了当前正在运行的所有事务的锁信息。该表位于信息模式（information_schema）中，其表结构如下所示。该表包含两个主键：requesting_trx_id 和 requesting_lock_id。该表与 INNO_DB_LOCK_WAITS 表一起使用，用于跟踪事务的锁等待情况。该表在 MySQL 4.0 版本中引入，并在 6.7 版本中进行了修改。

表 6-7 表 INNO_DB_LOCK_WAITS 的结构

| 字段 | 说明 | 字段 | 说明 |
|--------------------|-------------|------------------|----------|
| requesting_trx_id | 申请锁资源的事务 ID | blocking_trx_id | 阻塞的事务 ID |
| requesting_lock_id | 申请的锁的 ID | blocking_lock_id | 阻塞的锁的 ID |

表 INNO_DB_LOCK_WAITS 是 MySQL 数据库系统内部使用的一个表，它记录了当前正在运行的所有事务的锁等待信息。该表位于信息模式（information_schema）中，其表结构如下所示。该表包含两个主键：requesting_trx_id 和 requesting_lock_id。该表与 INNO_DB_LOCKS 表一起使用，用于跟踪事务的锁等待情况。该表在 MySQL 4.0 版本中引入，并在 6.7 版本中进行了修改。

```
mysql>SELECT*FROM information_schema.INNO_DB_LOCK_WAITS\G;

*****1.row*****

requesting_trx_id:7311F4

requested_lock_id:7311F4:96:3:2

blocking_trx_id:730FEE

blocking_lock_id:730FEE:96:3:2

1 row in set(0.00 sec)
```

表 INNO_DB_LOCK_WAITS 是 MySQL 数据库系统内部使用的一个表，它记录了当前正在运行的所有事务的锁等待信息。该表位于信息模式（information_schema）中，其表结构如下所示。该表包含两个主键：requesting_trx_id 和 requesting_lock_id。该表与 INNO_DB_LOCKS 表一起使用，用于跟踪事务的锁等待情况。该表在 MySQL 4.0 版本中引入，并在 6.7 版本中进行了修改。

```
mysql>SELECT

r.trx_id waiting_trx_id,
```

```
r.trx_mysql_thread_id waiting_thread,

r.trx_query waiting_query,

b.trx_id blocking_trx_id,

b.trx_mysql_thread_id blocking_thread,

b.trx_query blocking_query

FROM information_schema.innodb_lock_waits w

INNER JOIN information_schema.innodb_trx b

ON b.trx_id=w.blocking_trx_id

INNER JOIN information_schema.innodb_trx r

ON r.trx_id=w.requesting_trx_id\G;

*****1.row*****

waiting_trx_id:73122F

waiting_thread:471719

waiting_query:NULL

blocking_trx_id:7311FC

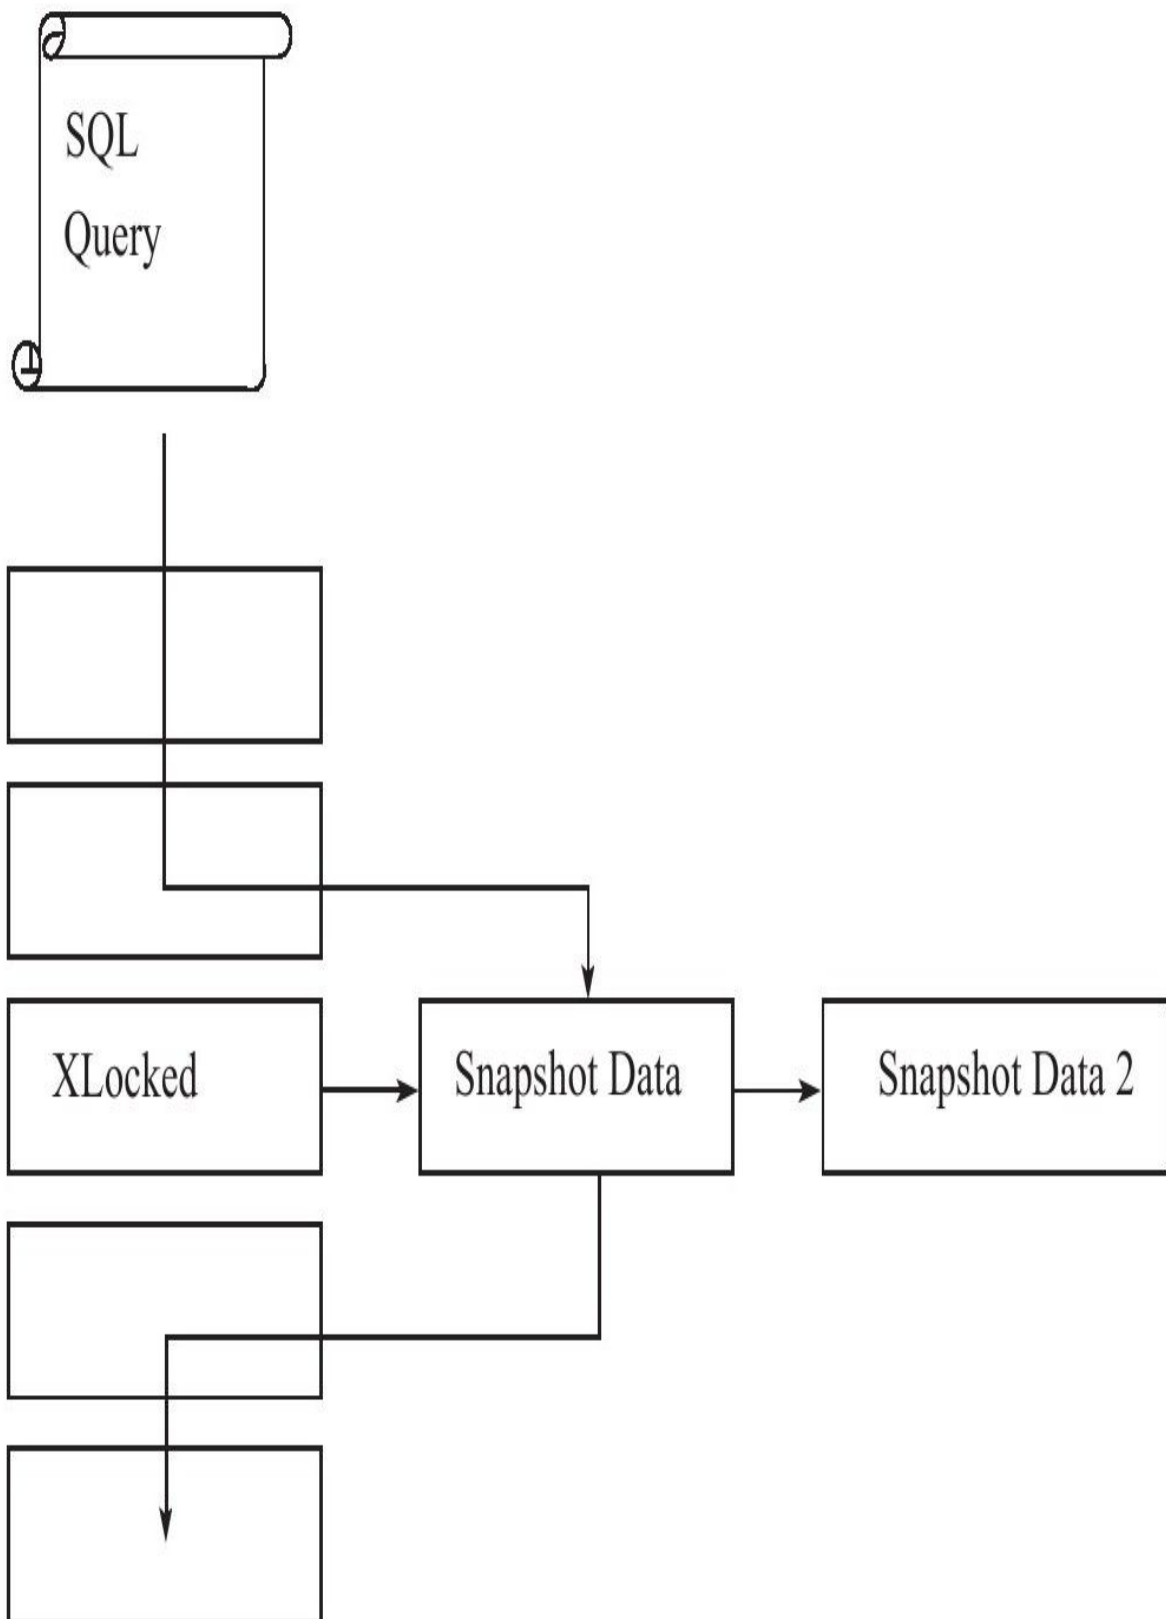
blocking_thread:471718

blocking_query:NULL

1 row in set(0.00 sec)
```

6.3.2 一致非锁定读

一致非锁定读 (consistent nonlocking read) 是 InnoDB 数据库的一个重要特性。它通过多版本控制 (multi versioning) 来实现。在 InnoDB 数据库中，DELETE 和 UPDATE 操作实际上是创建新的数据版本，而不是删除或修改现有版本。因此，即使存在并发操作，读者也可以看到一致的数据视图，而无需等待写操作完成。这一特性在 InnoDB 数据库的 6.4 版本中得到了进一步的优化。



6-4 InnoDB

6-4 InnoDB undo

InnoDB

6-4 Multi Version Concurrency Control MVCC

READ COMMITTED REPEATABLE READ InnoDB
 InnoDB
 READ COMMITTED
 REPEATABLE READ
 MySQL A SQL

```
#Session A

mysql> BEGIN;

Query OK, 0 rows affected (0.00 sec)

mysql> SELECT * FROM parent WHERE id=1;

+-----+
|id|
+-----+

|1|

+-----+

1 row in set (0.00 sec)
```

```

A
BEGIN
parent_id 1
B
B

```

```
mysql>BEGIN;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>UPDATE parent SET id=3 WHERE id=1;
```

```
Query OK,1 row affected(0.00 sec)
```

```
Rows matched:1 Changed:1 Warnings:0
```

Session B updates parent id 1 to id=3
Session X attempts to read A id 1
InnoDB reads A id 1 with READ COMMITTED REPEATABLE READ
A id 1 is updated to id=3
SQL: SELECT * FROM parent WHERE id=1
InnoDB reads A id 1 with READ COMMITTED REPEATABLE READ
A id 1 is updated to id=3

```
mysql>SELECT * FROM parent WHERE id=1;
```

```
+----+
```

```
|id|
```

```
+----+
```

```
|1|
```

```
+----+
```

```
1 row in set(0.00 sec)
```

Session id=1 updates id 1 to id=3
Session B reads A id 1

```
#Session B
```

```
mysql>commit;
```

```
Query OK,0 rows affected(0.01 sec)
```

Session B reads A id 1
SQL: SELECT * FROM parent WHERE id=1
InnoDB reads A id 1 with READ COMMITTED REPEATABLE READ
A id 1 is updated to id=3
InnoDB reads A id 1 with READ COMMITTED REPEATABLE READ
A id 1 is updated to id=3

fresh snapshot B READ COMMITTED

```
mysql> SELECT @@tx_isolation\G;
```

```
*****1.row*****
```

```
@@tx_isolation:READ-COMMITTED
```

```
1 row in set (0.00 sec)
```

```
mysql> SELECT * FROM parent WHERE id=1;
```

```
Empty set (0.00 sec)
```

REPEATABLE READ

```
mysql> SELECT @@tx_isolation\G;
```

```
*****1.row*****
```

```
@@tx_isolation:REPEATABLE-READ
```

```
1 row in set (0.00 sec)
```

```
mysql> SELECT * FROM parent WHERE id=1;
```

```
+----+
```

```
|id|
```

```
+----+
```

```
|1|
```

```
+----+
```

```
1 row in set (0.00 sec)
```

6-8 READ COMMITTED ACID 7

表 6-8 示例执行的过程

| 时间 | 会话 A | 会话 B |
|----|---------------------------------------|---|
| 1 | BEGIN | |
| 2 | SELECT * FROM parent
WHERE id = 1; | |
| 3 | | BEGIN |
| 4 | | UPDATE parent SET id=3
WHERE id = 1; |
| 5 | SELECT * FROM parent
WHERE id = 1; | |
| 6 | | COMMIT; |
| 7 | SELECT * FROM parent
WHERE id = 1; | |
| 8 | COMMIT | |

6.3.3 REPEATABLE READ

REPEATABLE READ InnoDB
SELECT
SELECT InnoDB
locking read

SELECT...FOR UPDATE

SELECT...LOCK IN SHARE MODE

SELECT...FOR UPDATE X
SELECT...LOCK IN SHARE MODE S
S X

SELECT...FOR UPDATE
SELECT...FOR UPDATE
SELECT...LOCK IN
SHARE MODE
SELECT
BEGIN
START TRANSACTION
SET
AUTOCOMMIT=0

6.3.4 自增列

MySQL 5.1.22 版本开始，InnoDB 支持 auto-increment counter 列。该列在插入新记录时，会自动生成一个唯一的值。

```
SELECT MAX(auto_inc_col) FROM t FOR UPDATE;
```

1. AUTO-INC Locking
MySQL 5.1.22 版本开始，InnoDB 支持 auto-increment counter 列。该列在插入新记录时，会自动生成一个唯一的值。

AUTO-INC Locking
INSERT...SELECT

MySQL 5.1.22 版本开始，InnoDB 支持 auto-increment counter 列。该列在插入新记录时，会自动生成一个唯一的值。
innodb_autoinc_lock_mode 1
6-9

表 6-9 插入类型

| 插入类型 | 说明 |
|--------------------|---|
| insert-like | insert-like 指所有的插入语句，如 INSERT、REPLACE、INSERT...SELECT、REPLACE...SEECT、LOAD DATA 等 |
| simple inserts | simple inserts 指能在插入前就确定插入行数的语句。这些语句包括 INSERT、REPLACE 等。需要注意的是：simple inserts 不包含 INSERT ...ON DUPLICATE KEY UPDATE 这类 SQL 语句 |
| bulk inserts | bulk inserts 指在插入前不能确定得到插入行数的语句，如 INSERT...SELECT、REPLACE...SELECT、LOAD DATA |
| mixed-mode inserts | mixed-mode inserts 指插入中有一部分的值是自增长的，有一部分是确定的。如 INSERT INTO t1 (c1,c2) VALUES (1,'a'), (NULL,'b'), (5,'c'), (NULL,'d'); 也可以是指 INSERT ...ON DUPLICATE KEY UPDATE 这类 SQL 语句 |

innodb_autoinc_lock_mode 0 1 2 6-10

表 6-10 参数 innodb_autoinc_lock_mode 的说明

| innodb_autoinc_lock_mode | 说明 |
|--------------------------|--|
| 0 | 这是 MySQL5.1.22 版本之前自增长的实现方式，即通过表锁的 AUTO-INC Locking 方式。因为有了新的自增长实现方式，0 这个选项不应该是新版用户的首选选项 |
| 1 | 这是该参数的默认值。对于“simple inserts”，该值会用互斥量（mutex）去对内存中的计数器进行累加的操作。对于“bulk inserts”，还是使用传统表锁的 AUTO-INC Locking 方式。在这种配置下，如果不考虑回滚操作，对于自增值列的增长还是连续的。并且在这种方式下，statement-based 方式的 replication 还是能很好地工作。需要注意的是，如果已经使用 AUTO-INC Locing 方式去产生自增长的值，而这时需要再进行“simple inserts”的操作时，还是需要等待 AUTO-INC Locking 的释放 |
| 2 | 在这个模式下，对于所有“INSERT-like”自增长值的产生都是通过互斥量，而不是 AUTO-INC Locking 的方式。显然，这是性能最高的方式。然而，这会带来一定的问题。因为并发插入的存在，在每次插入时，自增长的值可能不是连续的。此外，最重要的是，基于 Statement-Base Replication 会出现问题。因此，使用这个模式，任何时候都应该使用 row-base replication。这样才能保证最大的并发性能及 replication 主从数据的一致 |

MySQL 5.6 版本开始支持 InnoDB 引擎的 MyISAM 表复制。
MySQL 5.6 版本开始支持 InnoDB 引擎的 master-slave
MyISAM 引擎的 replication 功能。

MySQL 5.6 版本开始支持 InnoDB 引擎的 MyISAM 表复制。
MySQL 5.6 版本开始支持 InnoDB 引擎的 master-slave
MyISAM 引擎的 replication 功能。

```
mysql> CREATE TABLE t(
```

```
- a INT AUTO_INCREMENT,
```

```
- b INT,
```

```
- KEY(b,a)
```

```
- )ENGINE=InnoDB;
```

ERROR 1075 (42000): Incorrect table definition; there can be only one auto column and it must be defined as a key

```
mysql> CREATE TABLE t(
```

```
- a INT AUTO_INCREMENT,
```

```
- b INT,
```

```
- KEY(b,a)
```

```
- )ENGINE=MyISAM;
```

Query OK, 0 rows affected (0.01 sec)

6.3.5 外键

Oracle 数据库默认使用 InnoDB 存储引擎。InnoDB 存储引擎支持外键。Oracle 数据库默认使用 Oracle 存储引擎。Oracle 存储引擎不支持外键。

在 Oracle 数据库中，使用 SELECT 语句时，可以使用 LOCK IN SHARE MODE 选项。在 InnoDB 数据库中，使用 SELECT 语句时，可以使用 SELECT ... LOCK IN SHARE MODE 选项。在 Oracle 数据库中，使用 SELECT 语句时，可以使用 SELECT ... LOCK IN SHARE MODE 选项。

表 6-11 外键测试用例

| 时 间 | 会话 A | 会话 B |
|-----|--------------------------------|--|
| 1 | BEGIN | |
| 2 | DELETE FROM parent WHERE id=3; | |
| 3 | | BEGIN |
| 4 | | INSERT INTO child SELECT 2,3
第二列是外键，执行该句时被阻塞
(waiting) |

在 Oracle 数据库中，使用 COMMIT 和 ROLLBACK 语句。在 InnoDB 数据库中，使用 COMMIT 和 ROLLBACK 语句。在 Oracle 数据库中，使用 COMMIT 和 ROLLBACK 语句。

Session B id=3 A
id=3 INNODB_LOCKS

```
mysql>SELECT*FROM information_schema.INNODB_LOCKS\G;
```

```
*****1.row*****
```

```
lock_id:7573B8:96:3:4
```

```
lock_trx_id:7573B8
```

```
lock_mode:S
```

```
lock_type:RECORD
```

```
lock_table:'mytest'.'parent'
```

```
lock_index:'PRIMARY'
```

```
lock_space:96
```

```
lock_page:3
```

```
lock_rec:4
```

```
lock_data:3
```

```
*****2.row*****
```

```
lock_id:7573B3:96:3:4
```

```
lock_trx_id:7573B3
```

```
lock_mode:X
```

```
lock_type:RECORD
```

```
lock_table:'mytest'.'parent'
```

```
lock_index:'PRIMARY'
```

```
lock_space:96
```

```
lock_page:3
```

```
lock_rec:4
```

```
lock_data:3
```

```
2 rows in set(0.00 sec)
```

6.4 数据库

6.4.1 数据库3种锁

InnoDB数据库3种锁分别是

□ Record Lock 记录锁

□ Gap Lock 间隙锁

□ Next-Key Lock: Gap Lock + Record Lock 下一键锁

Record Lock 记录锁 InnoDB 数据库在插入、更新、删除操作时，会对记录加锁。InnoDB 数据库在插入、更新、删除操作时，会对记录加锁。

Next-Key Lock 下一键锁 Gap Lock 间隙锁 Record Lock 记录锁 Next-Key Lock 下一键锁 InnoDB 数据库在插入、更新、删除操作时，会对记录加锁。10 11 13 20 下一键锁 Next-Key Locking 下一键锁

$(-\infty, 10]$

$(10, 11]$

$(11, 13]$

$(13, 20]$

$(20, +\infty)$

□ Next-Key Lock 下一键锁 Next-Key Locking 下一键锁 Phantom Problem 幻读问题 predict lock 预测锁 next-key locking 下一键锁 previous-key locking 上一键锁 10 11 13 20 上一键锁

$(-\infty, 10)$

$[10, 11)$

表 6-12 唯一索引的锁定示例

| 时 间 | 会话 A | 会话 B |
|-----|--|-------------------------|
| 1 | BEGIN; | |
| 2 | SELECT * FROM t
WHERE a=5 FOR UPDATE; | |
| 3 | | BEGIN; |
| 4 | | INSERT INTO t SELECT 4; |
| 5 | | COMMIT;
成功, 不需要等待 |
| 6 | COMMIT | |

t 1 2 5 唯一索引的锁定示例 A 的 a=5 的 X 锁阻止 B 的 a=5 的插入 (2 5) 的插入 B 的 4 的插入成功, 不需要等待
 Next-Key Lock 阻止 Record Lock 阻止插入

Next-Key Lock 阻止 Record Lock 阻止插入
 阻止插入 z

```
CREATE TABLE z(a INT,b INT,PRIMARY KEY(a),KEY(b));
```

```
INSERT INTO z SELECT 1,1;
```

```
INSERT INTO z SELECT 3,1;
```

```
INSERT INTO z SELECT 5,3;
```

```
INSERT INTO z SELECT 7,6;
```

```
INSERT INTO z SELECT 10,8;
```

z b A SQL

```
SELECT*FROM z WHERE b=3 FOR UPDATE
```

SQL b Next-Key Lock
a 5 Record Lock
Next-Key Lock (1 3) InnoDB
gap lock (3 6)
B SQL

```
SELECT*FROM z WHERE a=5 LOCK IN SHARE MODE;
```

```
INSERT INTO z SELECT 4,2;
```

```
INSERT INTO z SELECT 6,5;
```

SQL A SQL a=5 X
SQL 4 2 (1 3)
SQL 6 5 (1 3) 5
(3 6) SQL

```
INSERT INTO z SELECT 8,6;
```

```
INSERT INTO z SELECT 2,0;
```

```
INSERT INTO z SELECT 6,7;
```

Gap Lock
Phantom Problem A b=3

Gap Lock 3 6 b 3 A Phantom Problem

Gap Lock

READ COMMITTED

innodb_locks_unsafe_for_binlog 1

Gap Lock Record Lock replication READ COMMITTED READ REPEATABLE

InnoDB INSERT A z b=3 (1 3)

INSERT INTO z SELECT 2,2;

b 2 3

INSERT INTO z SELECT 2,0;

Next-Key Lock Record Lock range point InnoDB Next-Key Lock

6.4.2 Phantom Problem

REPEATABLE READ InnoDB Next-Key Locking Phantom Problem
Oracle SERIALIZABLE Phantom Problem

Phantom Problem SQL
SQL t t 1 2 5
T1 SQL

```
SELECT * FROM t WHERE a < 2 FOR UPDATE;
```

T1 5 T2 4
T1 SQL 4 5
6-13

表 6-13 Phantom Problem 的演示

| 时 间 | 会话 A | 会话 B |
|-----|--|------|
| 1 | SET SESSION
tx_isolation='READ-OMMITTED'; | |
| 2 | BEGIN; | |
| 3 | SELECT * FROM t
WHERE a > 2 FOR UPDATE;
***** 1. row *****

a: 4 | |

(续)

| 时 间 | 会话 A | 会话 B |
|-----|--|-------------------------|
| 4 | | BEGIN; |
| 5 | | INSERT INTO t SELECT 4; |
| 6 | | COMMIT; |
| 7 | SELECT * FROM t
WHERE a > 2 FOR UPDATE;
***** 1. row *****

a: 4

***** 2. row *****

a: 5 | |

InnoDB 使用 Next-Key Locking 防止 Phantom Problem
SQL: SELECT * FROM t WHERE a > 2 FOR UPDATE
InnoDB 使用 2- ∞ 范围锁防止 X 事务插入新行
Phantom Problem

InnoDB 支持 REPEATABLE READ
Next-Key Locking 防止 READ COMMITTED 幻读

Record Lock A READ COMMITTED

InnoDB Next-Key Locking

```
SELECT * FROM table WHERE col=xxx LOCK IN SHARE MODE
```

If not found any row:

#unique for insert value

```
INSERT INTO table VALUES(...);
```

S Lock SELECT... LOCK IN SHARE MODE 6-14

表 6-14 通过 Next-Key Locking 实现应用程序的唯一性检查

| 时 间 | 会话 A | 会话 B |
|-----|---|---|
| 1 | BEGIN | |
| 2 | mysql>SELECT * FROM z
WHERE b=4

LOCK IN SHARE MODE; | |
| 3 | | mysql>SELECT * FROM z
WHERE b=4

LOCK IN SHARE MODE; |
| 4 | mysql>INSERT INTO z SELECT 4,4;
阻塞 | |
| 5 | | mysql>INSERT INTO z
SELECT 4,4;

ERROR 1213 (40001):Deadlock found when
trying to get lock;try restarting transaction
抛出死锁异常 |
| 6 | # INSERT 插入成功 | |

6.5 脏读

脏读是指一个事务读取了另一个事务尚未提交的数据。在数据库系统中，脏读通常发生在未提交的事务被其他事务读取时。

6.5.1 脏读

脏读（Dirty Read）是指一个事务读取了另一个事务尚未提交的数据。在数据库系统中，脏读通常发生在未提交的事务被其他事务读取时。脏读会导致数据不一致，因此在数据库设计中需要避免脏读。

脏读通常发生在未提交的事务被其他事务读取时。在数据库系统中，脏读通常发生在未提交的事务被其他事务读取时。脏读会导致数据不一致，因此在数据库设计中需要避免脏读。

脏读通常发生在未提交的事务被其他事务读取时。在数据库系统中，脏读通常发生在未提交的事务被其他事务读取时。脏读会导致数据不一致，因此在数据库设计中需要避免脏读。

脏读通常发生在未提交的事务被其他事务读取时。在数据库系统中，脏读通常发生在未提交的事务被其他事务读取时。脏读会导致数据不一致，因此在数据库设计中需要避免脏读。6-15

表 6-15 脏读的示例

| Time | 会话 A | 会话 B |
|------|--|--|
| 1 | SET
@@tx_isolation='read-ncommitted'; | |
| 2 | | SET
@@tx_isolation='read-ncommitted'; |
| 3 | | BEGIN; |
| 4 | | mysql> SELECT * FROM t\G;
***** 1. row *****

a: 1

1 row in set (0.00 sec) |
| 5 | INSERT INTO t SELECT 2; | |
| 6 | | mysql> SELECT * FROM t\G;
***** 1. row *****

a: 1

***** 2. row *****

a: 2

2 row in set (0.00 sec) |

6.4.1 REPEATABLE READ READ UNCOMMITTED A
B SELECT 2 A
2

READ UNCOMMITTED READ
COMMITTED InnoDB READ REPEATABLE
Microsoft SQL Server READ COMMITTED Oracle
READ COMMITTED

READ
UNCOMMITTED replication slave slave

6.5.2 数据字典

数据字典是数据库系统的重要组成部分，它用于存储数据库的元数据。数据字典通常包含数据库的结构信息，如表名、列名、数据类型、约束条件等。数据字典是数据库管理系统（DBMS）的核心组件之一，它为数据库的查询、更新、删除等操作提供必要的元数据支持。数据字典的维护和管理是数据库管理员（DBA）的重要职责之一。

数据字典的维护和管理通常包括以下几个方面：1. 数据字典的初始化：在数据库创建时，需要初始化数据字典，包括创建数据字典表、定义数据字典结构等。2. 数据字典的更新：当数据库结构发生变化时，需要及时更新数据字典，以反映最新的数据库结构。3. 数据字典的备份和恢复：为了确保数据字典的安全性和可用性，需要定期对数据字典进行备份，并在需要时进行恢复。4. 数据字典的监控和审计：需要定期对数据字典进行监控和审计，以确保其完整性和准确性。

6-16

表 6-16 不可重复读的示例

| Time | 会话 A | 会话 B |
|------|---|--------------------------------------|
| 1 | SET @@tx_isolation='read-committed'; | |
| 2 | | SET @@tx_isolation='read-committed'; |
| 3 | BEGIN | BEGIN |
| 4 | mysql>SELECT * FROM t;
***** 1. row *****

a: 1

1 row in set (0.00 sec) | |
| 5 | | INSERT INTO t SELECT 2; |
| 6 | | COMMIT; |
| 7 | mysql>SELECT * FROM t;
***** 1. row *****

a: 1

***** 1. row *****

a: 2

2 row in set (0.00 sec) | |

6.5.3 〇〇〇〇

[illegible]

100T1000r00v100000T100000

$$2 \frac{v_1}{r_1} T_1 = 2 \frac{v_2}{r_2} T_2$$

300T1000

4□□□T2□□□

```

READ UNCOMMITTED
2T2rrT1

```

[illegible]

```
1[T1]User1
```

2T2User2

3User1

4User2

```

00000000User10000000“”0000000000“”0000000000
000000000000000010 00000000000000000000000000
009000000000000000000000000000000000000000001
0000000000000000000000999900000000009000000000
000000000000000000000090000000000000000000000
00000000USB Key0000000000000000USB Key0000000000
0000000000000000000000000000000000000000000

```

1
X2X
21346-17
1

表 6-17 丢失更新问题的处理方法

| Time | 会话 A | 会话 B |
|------|--|--|
| 1 | BEGIN; | |
| 2 | SELECT cash into @cash
FROM account
WHERE user = pUser FOR UPDATE; | |
| 3 | | SELECT cash into @cash
FROM account
WHERE user = pUser FOR UPDATE;
等待 |
| | | |
| m | UPDATE account
SET cash=@cash-9000
WHERE user=pUser | |
| m+1 | COMMIT | |
| m+2 | | UPDATE account SET cash=@cash-1
WHERE user=pUser; |
| m+3 | | COMMIT |

UPDATE...FOR
UPDATE...UPDATE...
UPDATE...
SELECT...UPDATE...SQL...

SELECT...INSERT...UPDATE...DELETE...
...
...

6.6 锁等待

MySQL 5.6 版本开始，InnoDB 锁等待超时时间（innodb_lock_wait_timeout）默认为 50 秒。如果超过这个时间，事务就会回滚（innodb_rollback_on_timeout 默认为 ON）。如果设置为 OFF，事务就会一直等待，直到锁被释放为止。

InnoDB 锁等待超时时间（innodb_lock_wait_timeout）默认为 50 秒。如果超过这个时间，事务就会回滚（innodb_rollback_on_timeout 默认为 ON）。如果设置为 OFF，事务就会一直等待，直到锁被释放为止。MySQL 5.6 版本开始，InnoDB 锁等待超时时间（innodb_lock_wait_timeout）默认为 50 秒。

```
mysql> SET @@innodb_lock_wait_timeout=60;
```

```
Query OK, 0 rows affected (0.00 sec)
```

innodb_rollback_on_timeout 默认为 ON。如果设置为 OFF，事务就会一直等待，直到锁被释放为止。

```
mysql> SET @@innodb_rollback_on_timeout=on;
```

```
ERROR 1238 (HY000): Variable 'innodb_rollback_on_timeout' is a read only variable
```

MySQL 5.6 版本开始，InnoDB 锁等待超时时间（innodb_lock_wait_timeout）默认为 50 秒。

```
mysql> BEGIN;
```

```
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> SELECT * FROM t WHERE a=1 FOR UPDATE;
```

```
ERROR 1205 (HY000): Lock wait timeout exceeded; try restarting transaction
```

InnoDB 锁等待超时时间（innodb_lock_wait_timeout）默认为 50 秒。如果超过这个时间，事务就会回滚（innodb_rollback_on_timeout 默认为 ON）。如果设置为 OFF，事务就会一直等待，直到锁被释放为止。

```
# A
```

```
mysql> SELECT * FROM t;
```

+---+

|a|

+---+

|1|

|2|

|4|

+---+

3 rows in set(0.00 sec)

mysql>BEGIN;

Query OK,0 rows affected(0.00 sec)

mysql>SELECT*FROM t WHERE a=4 FOR UPDATE;

+---+

|a|

+---+

|1|

|2|

+---+

2 rows in set(0.00 sec)

MySQL A Next-Key Lock 4
B

#B

mysql>BEGIN;

Query OK,0 rows affected(0.00 sec)

mysql>INSERT INTO t SELECT 5;

Query OK,1 row affected(0.00 sec)

Records:1 Duplicates:0 Warnings:0

mysql>INSERTINTO t SELECT 3;

ERROR 1205(HY000):Lock wait timeout exceeded;try restarting transaction

プロセスBが5秒間待たずに3秒間プロセスAのNext-Key Lockを
取得し、プロセスAが5秒間待たずにSELECT
を実行し、5秒間待たずに

```
mysql> SELECT * FROM t;
```

```
+---+
```

```
|a|
```

```
+---+
```

```
|1|
```

```
|2|
```

```
|4|
```

```
|5|
```

```
|8|
```

```
+---+
```

```
5 rows in set (0.00 sec)
```

プロセスBがCOMMITを実行し、
プロセスAがROLLBACKを実行し、
プロセスAがCOMMITを実行し、
プロセスAがROLLBACKを実行し、

6.7 □□

6.7.1 五五五五五

[illegible]

```

InnoDB innodb_lock_wait_timeout=30

```

```

000000000000000000000000000000000000FIFO00000000
000000000000000000000000000000000000undo log0000FIFO
000000000000000000000000000000000000

```

```

wait-for graph
InnoDB
wait-for
graph

```

□ □ □ □ □ □ □

□ □ □ □ □ □ □

```
wait-for graph T1 T2
```

□□T1□□□T2□□□□□□

☐ T1 ☐ T2

[illegible]

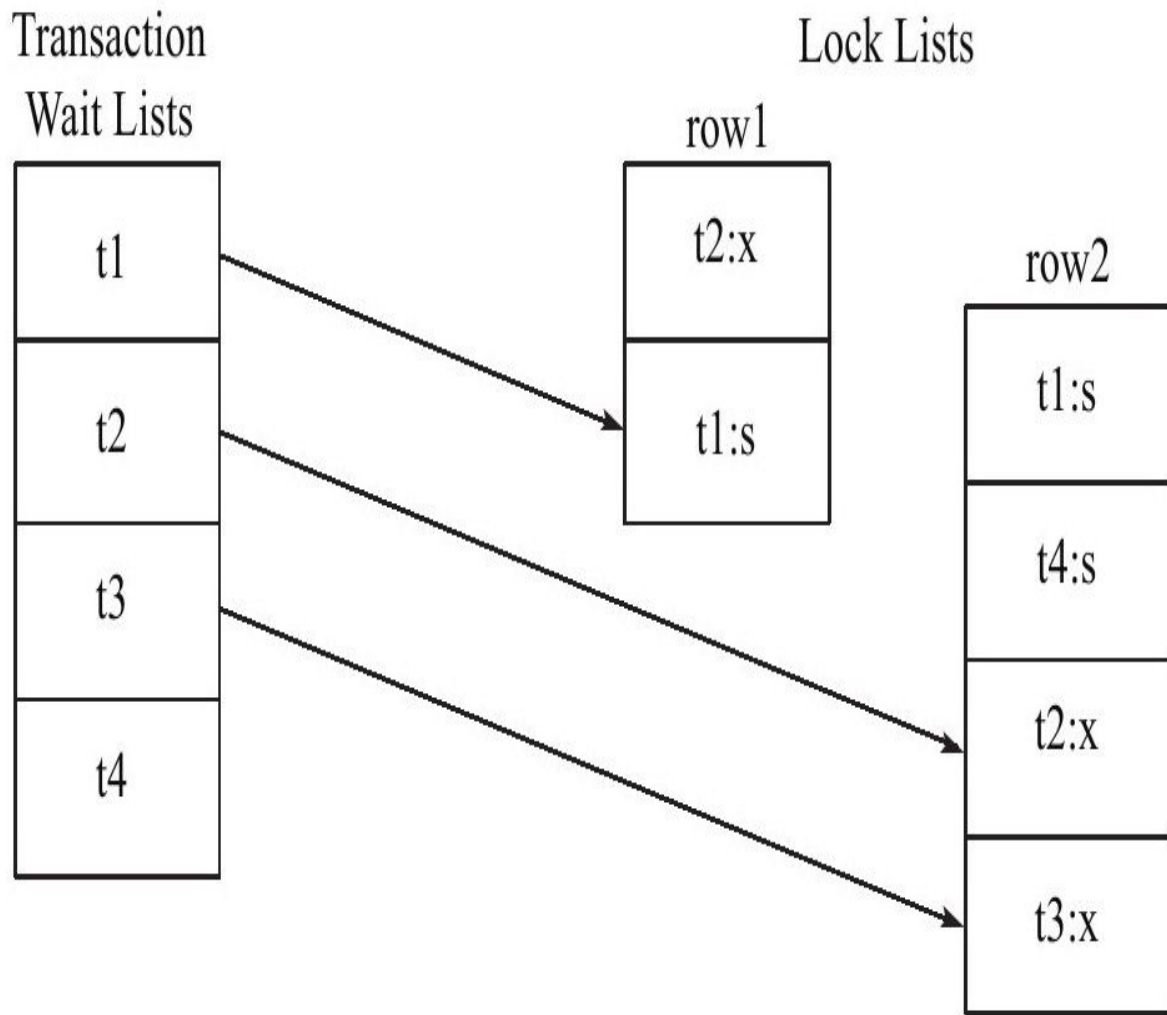


Figure 6-5 Transaction Wait Lists

Transaction Wait Lists: t1, t2, t3, t4
 Lock Lists:
 row1: t2:x, t1:s
 row2: t1:s, t4:s, t2:x, t3:x
 Mapping: t1 → row1 (t1:s), t2 → row2 (t2:x), t3 → row2 (t3:x), t4 → row2 (t4:s)

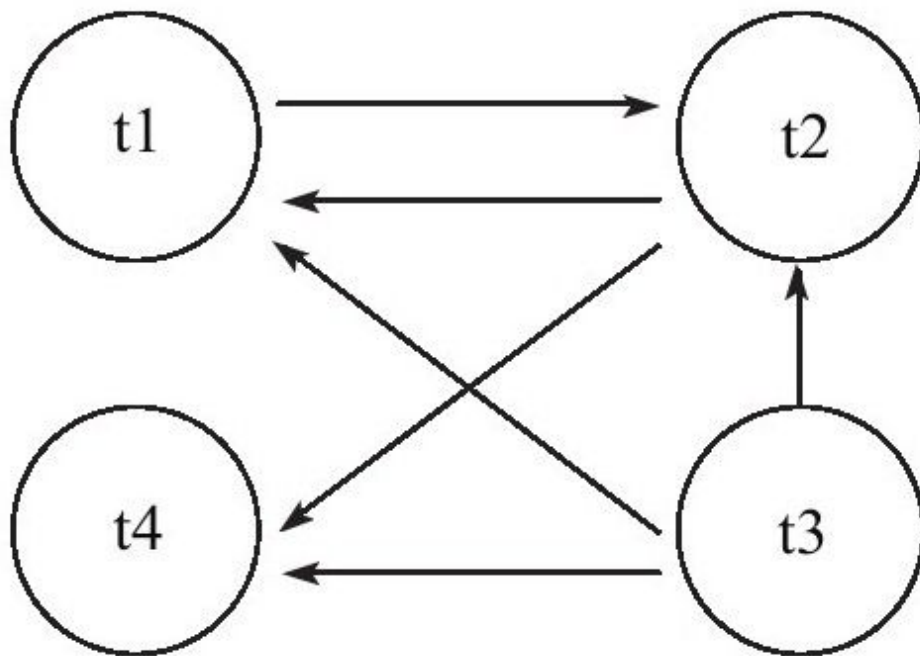


图 6-6 wait-for graph

图 6-6 展示了 InnoDB 的 wait-for graph。图中包含四个事务：t1、t2、t3 和 t4。t1 正在等待 t2 完成，t2 正在等待 t1 完成，t3 正在等待 t2 完成，t3 正在等待 t4 完成，t4 正在等待 t1 完成。InnoDB 的 undo 日志记录了这些事务的操作。

wait-for graph 是 InnoDB 1.2 版本引入的一个新功能。它用于检测死锁。InnoDB 1.2 版本在 wait-for graph 中引入了一个死锁检测算法，该算法可以检测出死锁并自动回滚其中一个事务。InnoDB 的 undo 日志记录了这些事务的操作。

6.7.2 死锁

考虑一个数据库系统，其中每个事务都包含一个或多个操作。每个操作都包含一个或多个数据项。每个数据项都有一个唯一的标识符。每个操作都包含一个或多个数据项的标识符。每个操作都包含一个或多个数据项的标识符。

假设每个事务都包含 $n+1$ 个操作。每个操作都包含 $n+1$ 个数据项。每个数据项都有一个唯一的标识符。每个操作都包含一个或多个数据项的标识符。每个操作都包含一个或多个数据项的标识符。

假设每个操作都包含 $n+1$ 个数据项。每个数据项都有一个唯一的标识符。每个操作都包含一个或多个数据项的标识符。每个操作都包含一个或多个数据项的标识符。

$$(1+2+3+\dots+r) / (r+1) \approx r/2$$

假设每个操作都包含 $n+1$ 个数据项。每个数据项都有一个唯一的标识符。每个操作都包含一个或多个数据项的标识符。每个操作都包含一个或多个数据项的标识符。

$$PW = nr/2R$$

假设每个操作都包含 $n+1$ 个数据项。每个数据项都有一个唯一的标识符。每个操作都包含一个或多个数据项的标识符。每个操作都包含一个或多个数据项的标识符。

$$PW(T) = 1 - (1 - PW)^r \approx r * PW \approx \frac{nr^2}{2R}$$

假设每个操作都包含 $n+1$ 个数据项。每个数据项都有一个唯一的标识符。每个操作都包含一个或多个数据项的标识符。每个操作都包含一个或多个数据项的标识符。

$$\text{一个事务发生死锁的概率} \approx \frac{PW(T)^2}{n} \approx \frac{nr^4}{4R^2}$$

假设每个操作都包含 $n+1$ 个数据项。每个数据项都有一个唯一的标识符。每个操作都包含一个或多个数据项的标识符。每个操作都包含一个或多个数据项的标识符。

系统中任何一个事务发生死锁的概率 $\approx \frac{n^2 r^4}{4R^2}$

nrR

[illegible]

□ □ □ □ □ □ □ □ r □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □

□ □ □ □ □ □ □ □ R □ □ □ □ □ □ □ □ □ □ □ □ □ □

6.7.3 同态

同态映射是代数系统之间的一种映射，它保持了代数系统的运算结构。设 $\langle A, \circ \rangle$ 和 $\langle B, \odot \rangle$ 是两个代数系统， $f: A \rightarrow B$ 是一个映射，如果对于任意的 $a, b \in A$ ，都有 $f(a \circ b) = f(a) \odot f(b)$ ，则称 f 是从 $\langle A, \circ \rangle$ 到 $\langle B, \odot \rangle$ 的同态映射。如果 f 还是单射和满射，则称 f 是同构映射。

例 6-18 设 $\langle A, \circ \rangle$ 和 $\langle B, \odot \rangle$ 是两个代数系统， $f: A \rightarrow B$ 是一个映射，如果对于任意的 $a, b \in A$ ，都有 $f(a \circ b) = f(a) \odot f(b)$ ，则称 f 是从 $\langle A, \circ \rangle$ 到 $\langle B, \odot \rangle$ 的同态映射。如果 f 还是单射和满射，则称 f 是同构映射。

AB-BA 同态

表 6-18 死锁用例 1

| 时间 | 会话 A | 会话 B |
|----|---|---|
| 1 | BEGIN; | |
| 2 | mysql>SELECT * FROM t
WHERE a = 1 FOR UPDATE;
***** 1. row *****
a: 1
1 row in set (0.00 sec) | BEGIN |
| 3 | | mysql>SELECT * FROM t
WHERE a = 2 FOR UPDATE;
***** 1. row *****
a: 2
1 row in set (0.00 sec) |
| 4 | mysql>SELECT * FROM t WHERE a = 2
FOR UPDATE;
等待 | |
| 5 | | mysql>SELECT * FROM t WHERE a = 1
FOR UPDATE;
ERROR 1213 (40001): Deadlock found when
trying to get lock; try restarting transaction |

Oracle 11gR2 1213 安装过程中，在安装 InnoDB 引擎时，会提示如下信息：

ORA-01111: 数据文件 2 的日志序列 6.6 与 InnoDB 引擎不兼容

解决方法：在安装 InnoDB 引擎时，将日志序列号改为 1213。

Oracle 11gR2 1213 安装过程中，在安装 InnoDB 引擎时，会提示如下信息：

ORA-01111: 数据文件 2 的日志序列 6.6 与 InnoDB 引擎不兼容

解决方法：在安装 InnoDB 引擎时，将日志序列号改为 1213。

```
mysql>CREATE TABLE p(  
-aINT,  
-PRIMARY KEY(a)  
-ENGINE=InnoDB;
```

Query OK, 0 rows affected (0.00 sec)

```
mysql>CREATE TABLE c(  
-bINT,  
-FOREIGN KEY(b)REFERENCES p(a)  
-ENGINE=InnoDB;
```

Query OK, 0 rows affected (0.00 sec)

```
mysql>SHOW INDEX FROM c\G;
```

*****1. row*****

Table: c

Non_unique: 1

Key_name: b

Seq_in_index: 1

Column_name: b

Collation: A

Cardinality: 0

Sub_part: NULL

Packed: NULL

Null:YES

Index_type:BTREE

Comment:

1 row in set(0.00 sec)

mysql>DROP INDEX b ON c;

ERROR 1553(HY000):Cannot drop index'b':needed in a foreign key constraint

MySQL InnoDB
b

MySQL X
S
t

CREATE TABLE t(

a INT PRIMARY KEY

)ENGINE=InnoDB;

INSERT INTO t VALUES(1),(2),(4),(5);

t a 4 6-19

表 6-19 死锁用例 2

| 时 间 | 会话 A | 会话 B |
|-----|--|--|
| 1 | BEGIN; | |
| 2 | | BEGIN; |
| 3 | SELECT * FROM t
WHERE a = 4 FOR UPDATE; | |
| 4 | | SELECT * FROM t
WHERE a <= 4 LOCK IN SHARE MODE;
-- 等待 |
| 5 | INSERT INTO t VALUES(3);
-- ERROR 1213 (40001): Deadlock found when
trying to get lock; try restarting transaction | |
| 6 | | -- 事务获得锁，正常运行 |

会话 A 在时间 4 请求 X 锁，会话 A 在时间 3 请求 S 锁，
 会话 B 在时间 4 请求 S 锁，会话 B 在时间 1 请求 S 锁，
 会话 B 在时间 4 请求 S 锁，会话 B 在时间 3 请求 S 锁。

InnoDB undo log AB-BA

6.8 問題

Lock Escalation 問題が発生する可能性がある場合、1000 以上のセッションが同時にデータベースに接続している場合、データベースは Lock Escalation を発生させる可能性がある。

Microsoft SQL Server は、Lock Escalation を発生させる可能性がある。Lock Escalation は、データベースが Lock を管理するために必要なリソースが不足している場合に発生する。

Microsoft SQL Server 2005 は、SQL Server の Lock Escalation を発生させる可能性がある。InnoDB は、Lock Escalation を発生させる可能性がある。

SQL は、Lock Escalation を発生させる可能性がある。5000 以上のセッションが同時にデータベースに接続している場合、データベースは Lock Escalation を発生させる可能性がある。

Lock Escalation は、データベースが Lock を管理するために必要なリソースが不足している場合に発生する。40% のセッションが同時にデータベースに接続している場合、データベースは Lock Escalation を発生させる可能性がある。

Microsoft SQL Server は、Lock Escalation を発生させる可能性がある。Lock Escalation は、データベースが Lock を管理するために必要なリソースが不足している場合に発生する。

InnoDB は、Lock Escalation を発生させる可能性がある。Lock Escalation は、データベースが Lock を管理するために必要なリソースが不足している場合に発生する。

3 000 000 以上のセッションが同時にデータベースに接続している場合、100 以上のセッションが同時にデータベースに接続している場合、300 000 000 以上のセッションが同時にデータベースに接続している場合、SQL は Lock Escalation を発生させる可能性がある。X は Lock Escalation を発生させる可能性がある。10 以上のセッションが同時にデータベースに接続している場合、3GB 以上のセッションが同時にデータベースに接続している場合、InnoDB は Lock Escalation を発生させる可能性がある。30 以上のセッションが同時にデータベースに接続している場合、90MB 以上のセッションが同時にデータベースに接続している場合、Lock Escalation は発生する可能性がある。

6.9 数据库

数据库是指长期存储在计算机内、有组织的、可共享的数据集合。数据库中的数据按一定的数据模型组织、存储于二级存储设备中，在用户应用程序中通过数据库管理系统进行数据管理。数据库系统是指数据库、数据库管理系统（DBMS）以及数据库管理员（DBA）的集合。数据库系统的主要功能包括：数据定义、数据操作、数据控制、数据查询、数据维护等。数据库系统的发展经历了多个阶段，从早期的文件系统到现在的关系数据库、NoSQL数据库等。数据库系统是现代信息系统的核心组成部分，广泛应用于各个领域。

数据库系统的主要类型包括：关系数据库（如MySQL、InnoDB、MyISAM）、非关系数据库（如Microsoft SQL Server、Oracle）、NoSQL数据库（如MongoDB、Cassandra）等。数据库系统的设计需要考虑多个因素，包括数据模型、数据完整性、数据安全、性能优化等。数据库系统的维护工作包括备份、恢复、性能监控、安全审计等。数据库系统是信息系统的基石，对于企业的数据管理和业务运营至关重要。

070 00

Transaction

[illegible]

InnoDB ACID ACID 4

atomicity

□□□□consistency□

□□□□isolation□

durability

6 InnoDB

7.1 四則演算

7.1.1 〇〇

SQL SQL

ACID
ACID
MySQL NDB Cluster D
Oracle READ COMMITTED I
InnoDB
READ REPEATABLE ACID
ACID

[illegible]

ATM

1 ATM

2□□□□□□□□□□□□□□□□

3 ATM

[illegible]

5 ATM

6□□□□□

ATM 数据库系统
数据库系统

SQL 数据库系统 SQL 数据库系统
数据库

数据库系统
数据库系统
数据库系统
数据库系统

C consistency 数据库系统
数据库系统
数据库系统
数据库系统
数据库系统——数据库系统

I isolation 数据库系统 concurrency control 数据库系统
serializability 数据库系统 locking 数据库系统
granular lock 数据库系统
数据库

D durability 数据库系统
数据库系统
数据库系统
RAID 数据库系统
High Reliability 数据库系统 High Availability 数据库系统

7.1.2 問題

□ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □

Flat Transactions

Flat Transactions with Savepoints

Chained Transactions

□□□□□Nested Transactions□

□□□□□□ Distributed Transactions □

```

Flat Transaction
BEGIN WORK COMMIT
WORK ROLLBACK WORK
7-1

```

BEGIN WORK
 Operation 1
 Operation 2

 Operation K
 COMMIT WORK

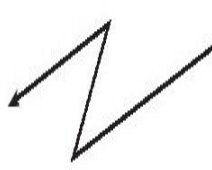
成功完成，约占所有事务的96%

BEGIN WORK
 Operation 1
 Operation 2

 (Error!!!)
 ROLLBACK WORK

应用程序要求停止事务，约占所有事务的3%

BEGIN WORK
 Operation 1
 Operation 2


 由于外界原因要回滚，如超时等

强制终止事务，约占所有事务的1%

7-1

7-1

BEGIN WORK

S1

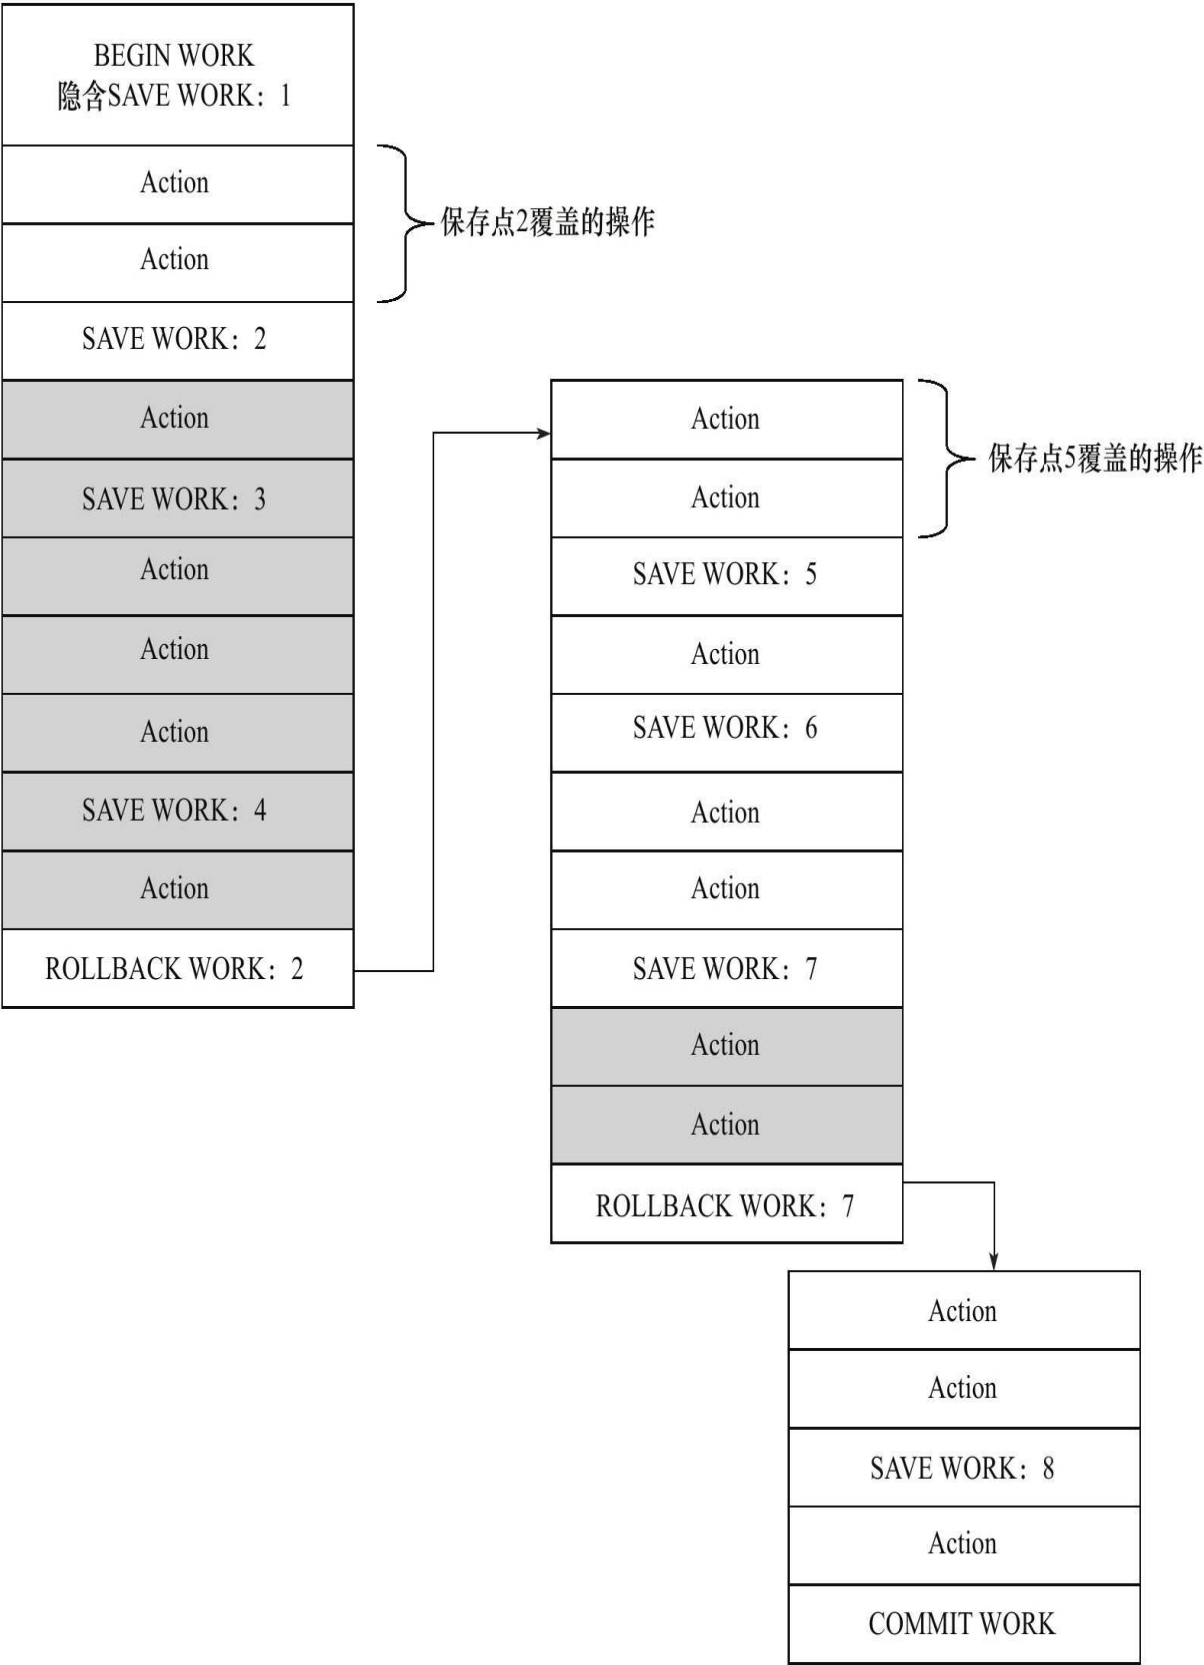
S2

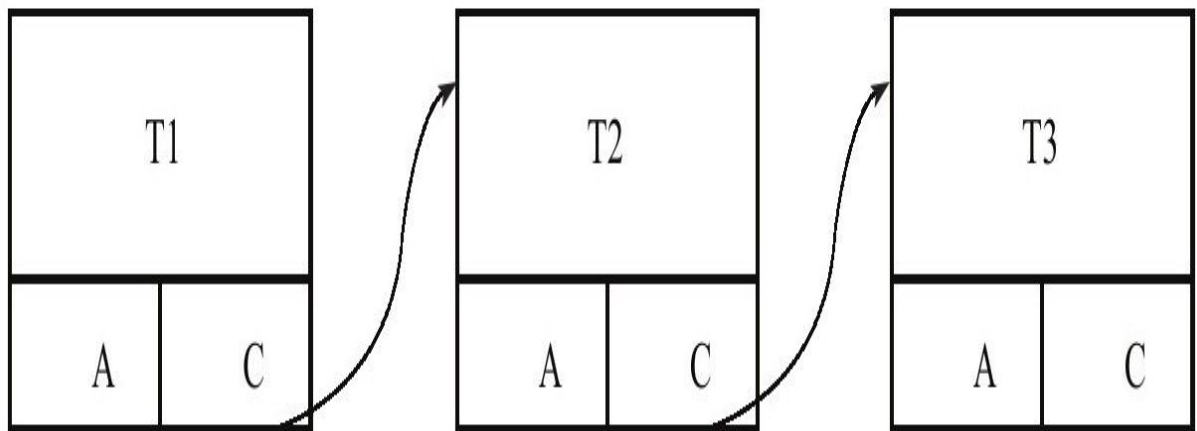
S3

S3
S1 S2 S3
S1 S2

Flat Transactions with Savepoint
Savepoint

SAVE WORK
7-2





触发器

触发器

图 7-3 数据库事务的触发器

数据库事务的触发器是指当数据库中的某个表发生数据变化时，系统会自动执行预先定义好的操作。这种操作可以是插入、更新或删除数据，也可以是调用存储过程或函数。触发器的主要作用是保证数据的一致性和完整性，防止非法数据的插入或修改。在数据库设计中，触发器通常用于实现复杂的业务逻辑，如审计、数据同步和备份等。当触发器被触发时，系统会自动执行触发器中定义的操作，直到遇到COMMIT语句为止。

数据库事务的触发器可以分为两种类型：一种是基于表的触发器，另一种是基于行的触发器。基于表的触发器是在整个表上定义的，当表中的数据发生变化时，触发器会被触发。基于行的触发器是在表的每一行上定义的，当表的某一行数据发生变化时，触发器会被触发。此外，数据库事务的触发器还可以分为两种类型：一种是基于事件的触发器，另一种是基于条件的触发器。基于事件的触发器是在某个事件发生时触发的，如插入、更新或删除数据。基于条件的触发器是在某个条件满足时触发的，如某个字段的值发生变化。

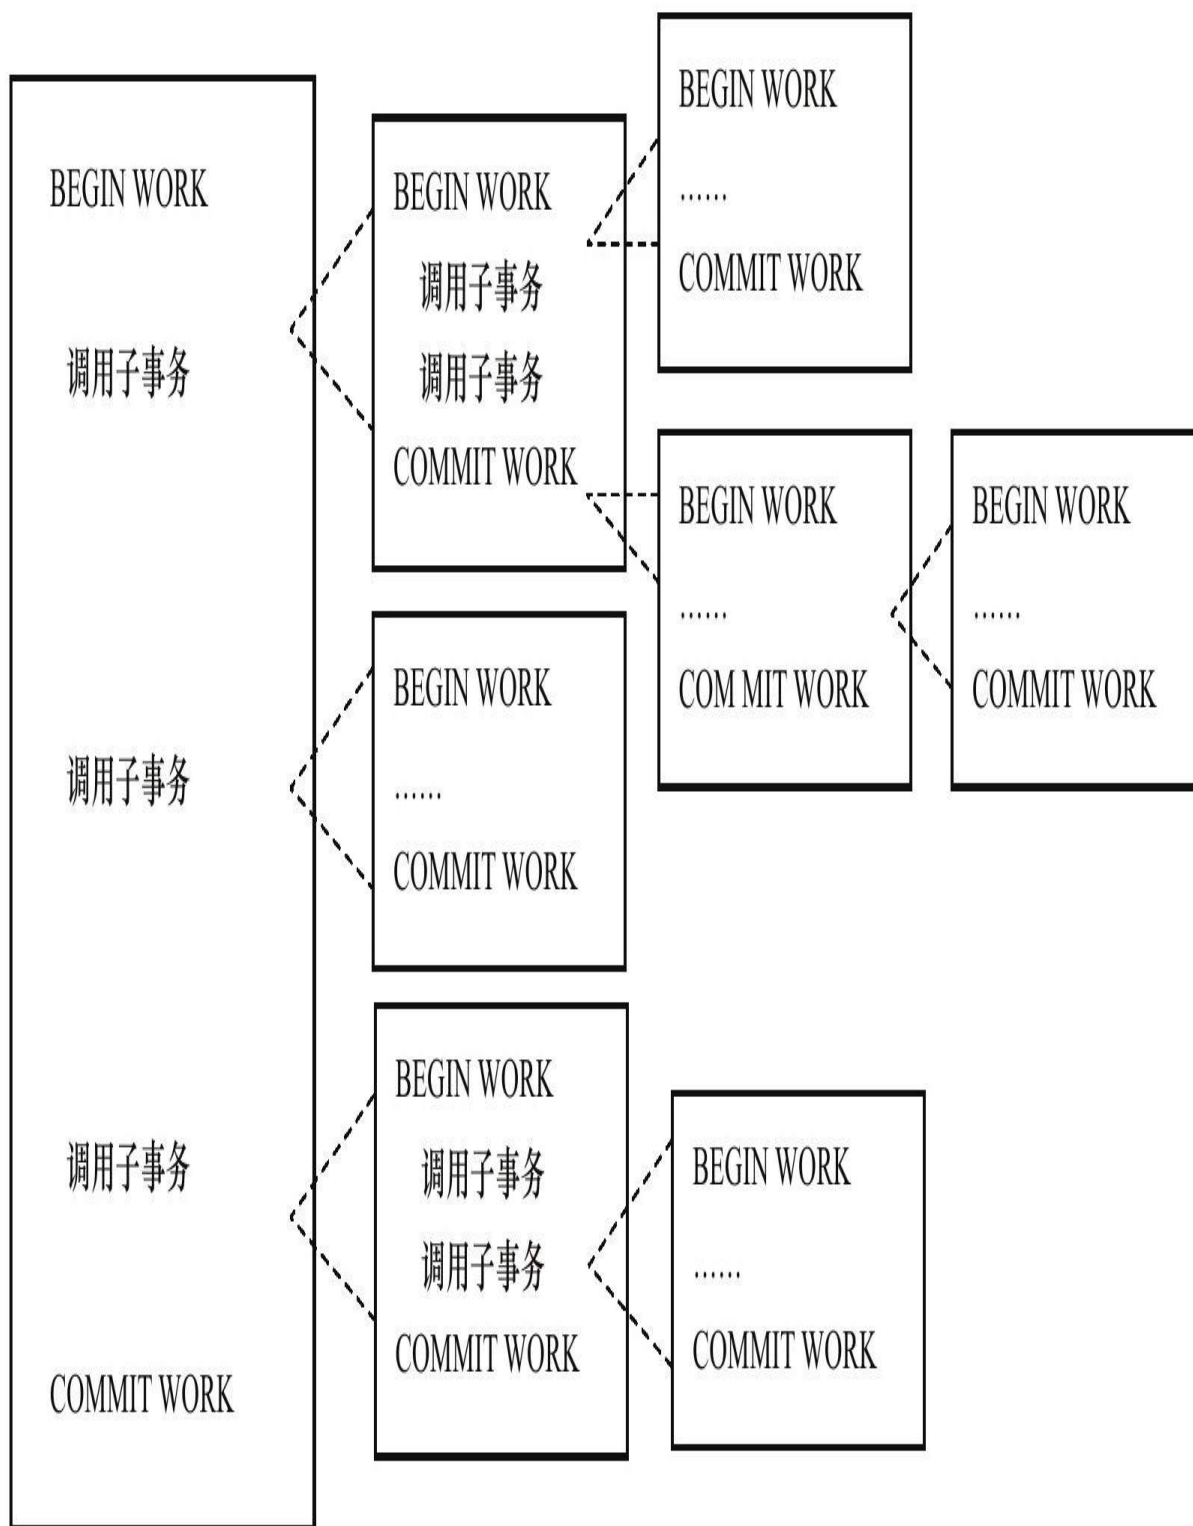
7-4

顶层事务

子事务

子事务

子事务



7-4 類別關係

Moss類別

1. 類別的定義

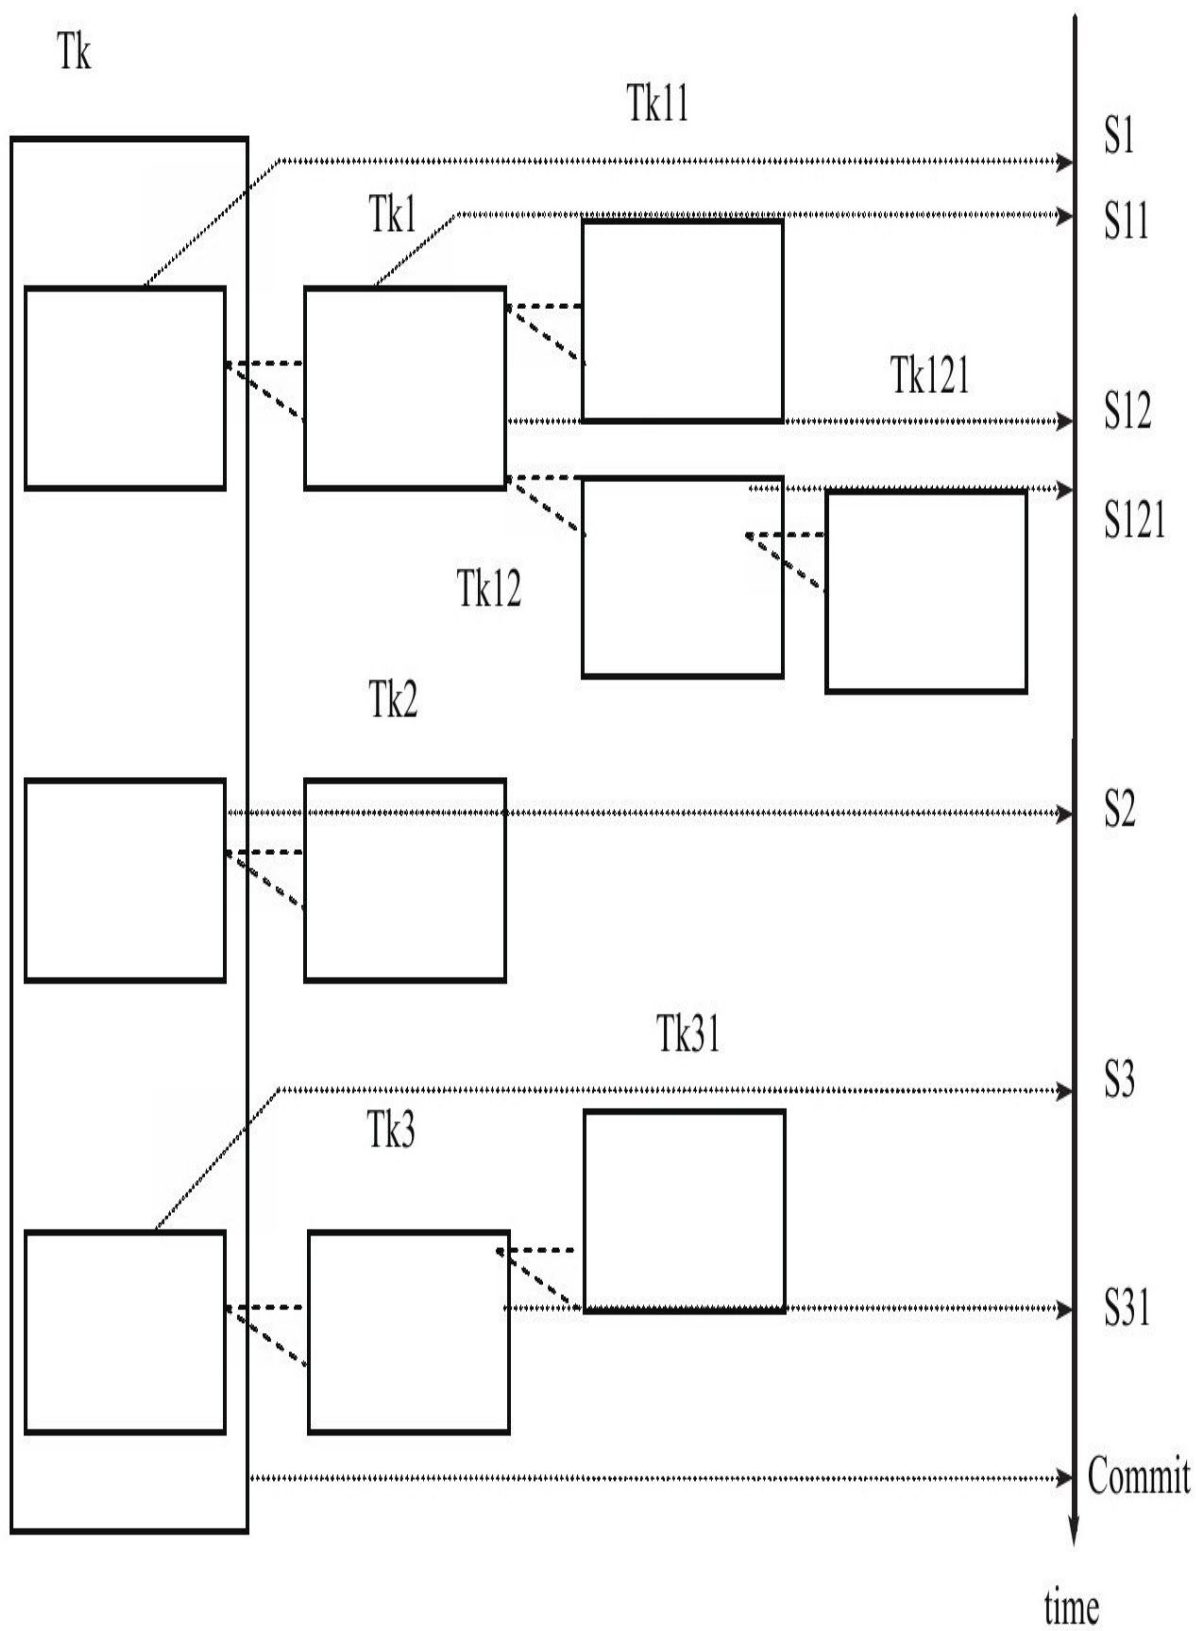
2. 類別的繼承

3. 類別的成員
predecessor
parent
child

4. 類別的初始化
初始化函數

5. 類別的構造函數
A C I D

Moss類別
類別的成員
類別的初始化
7-5



7-5

7-5 Tk3
S2

[illegible][illegible]

Distributed Transactions

[illegible]

1000A0000000

2000 B 10 000

[illegible]

4 A A

A

ACID

2 **3**

```

InnoDB
MySQL
InnoDB

```

7.2 五五五五五

```

000000060000000000000000000000000000redo logundo log00
00redo log000000000000000000000000undo log0000000000

```

```
DBAundoredoredoundo
redoundo
redoundo
```

7.2.1 redo

1.0000

```

ACID D redo
log buffer redo log file

```

```
InnoDB[ ]Force Log at Commit[ ]
[ ]COMMIT[ ]COMMIT[ ]
[ ]InnoDB[ ]redo log[ ]
undo log[ ]redo log[ ]undo log[ ]MVCC[ ]
redo log[ ]redo log[ ]undo
log[ ]
```

```

000000000000000000000000000000000000InnoDB00000000
0000fsync0000000000000000O_DIRECT0000000000000000
000000000000000000000000fsync0000fsync000000000000
00000000000000000000000000000000

```

[illegible]

```
innodb_flush_log_at_trx_commit 1
fsync 0
2
0
master thread master thread 1
```

fsync 2 MySQL
fsync MySQL
MySQL

innodb_flush_log_at_trx_commit
t1 p_load

```
CREATE TABLE test_load(  
  
  a INT,  
  
  b CHAR(80)  
  
)ENGINE=INNODB;  
  
DELIMITER//  
  
CREATE PROCEDURE p_load(count INT UNSIGNED)  
  
BEGIN  
  
  DECLARE s INT UNSIGNED DEFAULT 1;  
  
  DECLARE c CHAR(80)DEFAULT REPEAT('a',80);  
  
  WHILE s<=count DO  
  
    INSERT INTO test_load SELECT NULL,c;  
  
    COMMIT;  
  
    SET s=s+1;  
  
  END WHILE;  
  
  END;  
  
  //
```

p_load test_load
COMMIT innodb_flush_log_at_trx_commit 1
InnoDB fsync
CALL p_load 500 000 50 50 fsync
50

```
mysql>CALL p_load(500000);
```

```
Query OK,0 rows affected(1 min 53.11 sec)
```

```
innodb_flush_log_at_trx_commit=2
innodb_flush_log_at_trx_commit=0
```

```
mysql>SHOW VARIABLES LIKE'innodb_flush_log_at_trx_commit'\G
```

```
*****1.row*****
```

```
Variable_name:innodb_flush_log_at_trx_commit
```

```
Value:0
```

```
1 row in set(0.00 sec)
```

```
mysql>CALL p_load(500000);
```

```
Query OK,0 rows affected(13.90 sec)
```

```
innodb_flush_log_at_trx_commit=0
innodb_flush_log_at_trx_commit=50
innodb_flush_log_at_trx_commit=12%
innodb_flush_log_at_trx_commit=7-1
innodb_flush_log_at_trx_commit=p_load=50
```

表 7-1 不同 innodb_flush_log_at_trx_commit 设置对于插入的速度影响

| innodb_flush_log_at_trx_commit | 执行所用时间 |
|--------------------------------|-------------|
| 0 | 13.90 秒 |
| 1 | 1 分 53.11 秒 |
| 2 | 23.37 秒 |

innodb_flush_log_at_trx_commit=0/2 是 MySQL 数据库默认设置，它符合 ACID 特性，并且每秒 flush 一次。50 个事务提交后，MySQL 数据库会 flush 一次。COMMIT 操作会立即 flush。

MySQL 数据库支持 binlog 和 POINT-IN-TIME (PIT) 恢复。Replication 操作会立即 flush。

InnoDB 是 MySQL 数据库的默认存储引擎。MySQL 数据库支持 InnoDB 存储引擎。MySQL 数据库支持 InnoDB 存储引擎。

MySQL 数据库支持 InnoDB 存储引擎。MySQL 数据库支持 InnoDB 存储引擎。

MySQL 数据库支持 InnoDB 存储引擎。MySQL 数据库支持 InnoDB 存储引擎。

binlog

| | | | | | | | |
|----|----|----|----|----|----|----|----|
| T1 | T4 | T3 | T2 | T8 | T6 | T7 | T5 |
|----|----|----|----|----|----|----|----|

redo log

| | | | | | | | |
|----|----|----|-----|----|----|-----|-----|
| T1 | T2 | T1 | *T2 | T3 | T1 | *T3 | *T1 |
|----|----|----|-----|----|----|-----|-----|

图 7-6 事务日志的 redo log 和 binlog

图 7-6 展示了 InnoDB 事务日志的 redo log 和 binlog。redo log 记录了事务的 redo 操作，而 binlog 记录了事务的 redo 操作。图中展示了 redo log 和 binlog 的 redo log 和 binlog。图中展示了 redo log 和 binlog 的 redo log 和 binlog。

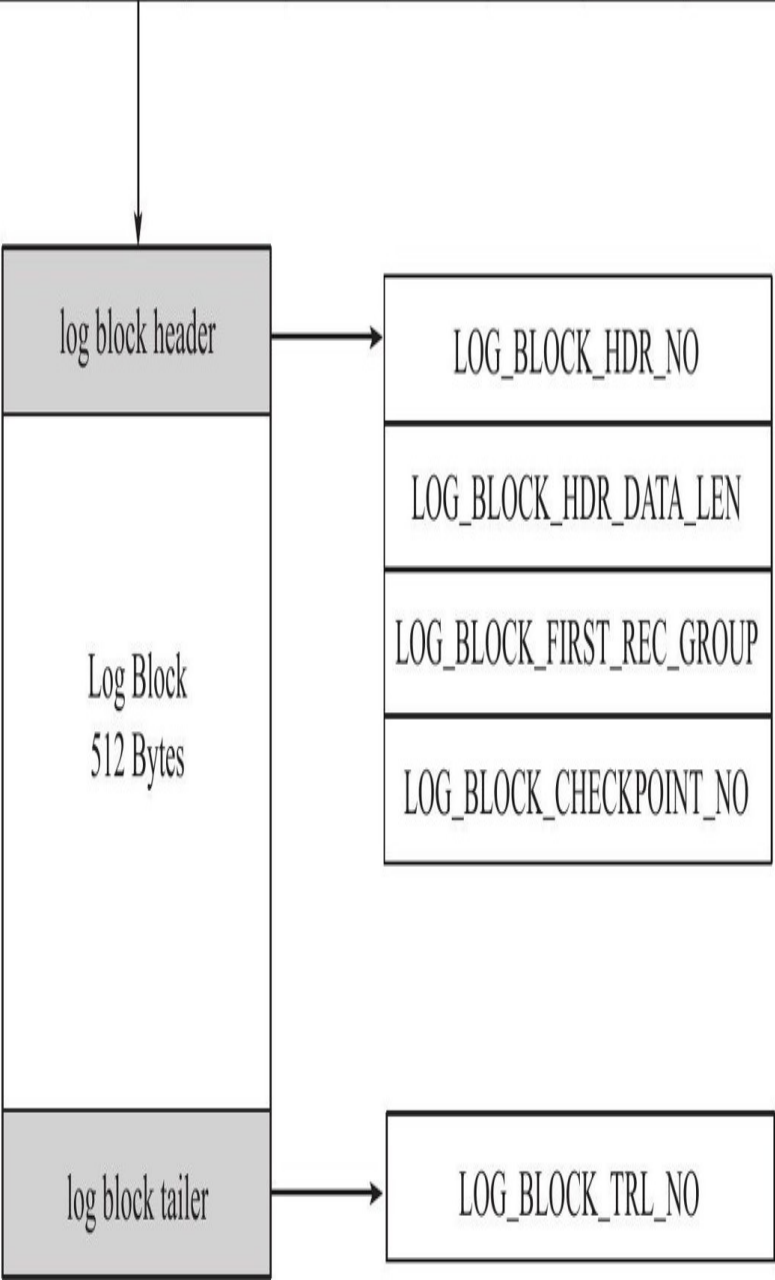
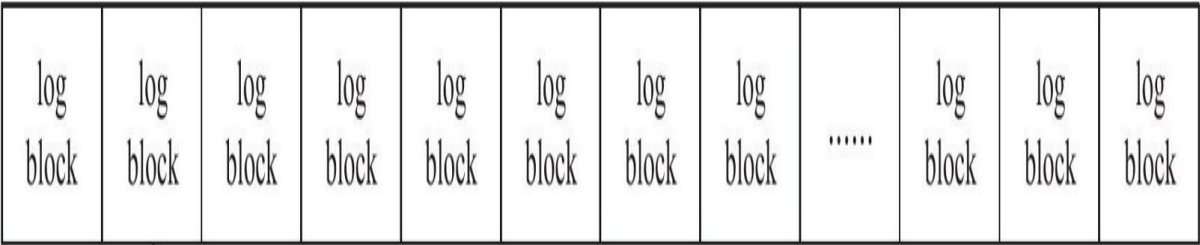
2.log block

InnoDB 事务日志的 redo log block 大小为 512 字节。redo log block 的大小为 512 字节。

redo log block 的大小为 512 字节。redo log block 的大小为 512 字节。redo log block 的大小为 512 字节。

redo log block header 的大小为 12 字节。redo log block header 的大小为 12 字节。redo log block header 的大小为 12 字节。

Redo Log Buffer



7-7

[illegible]

log block header 4 7-2

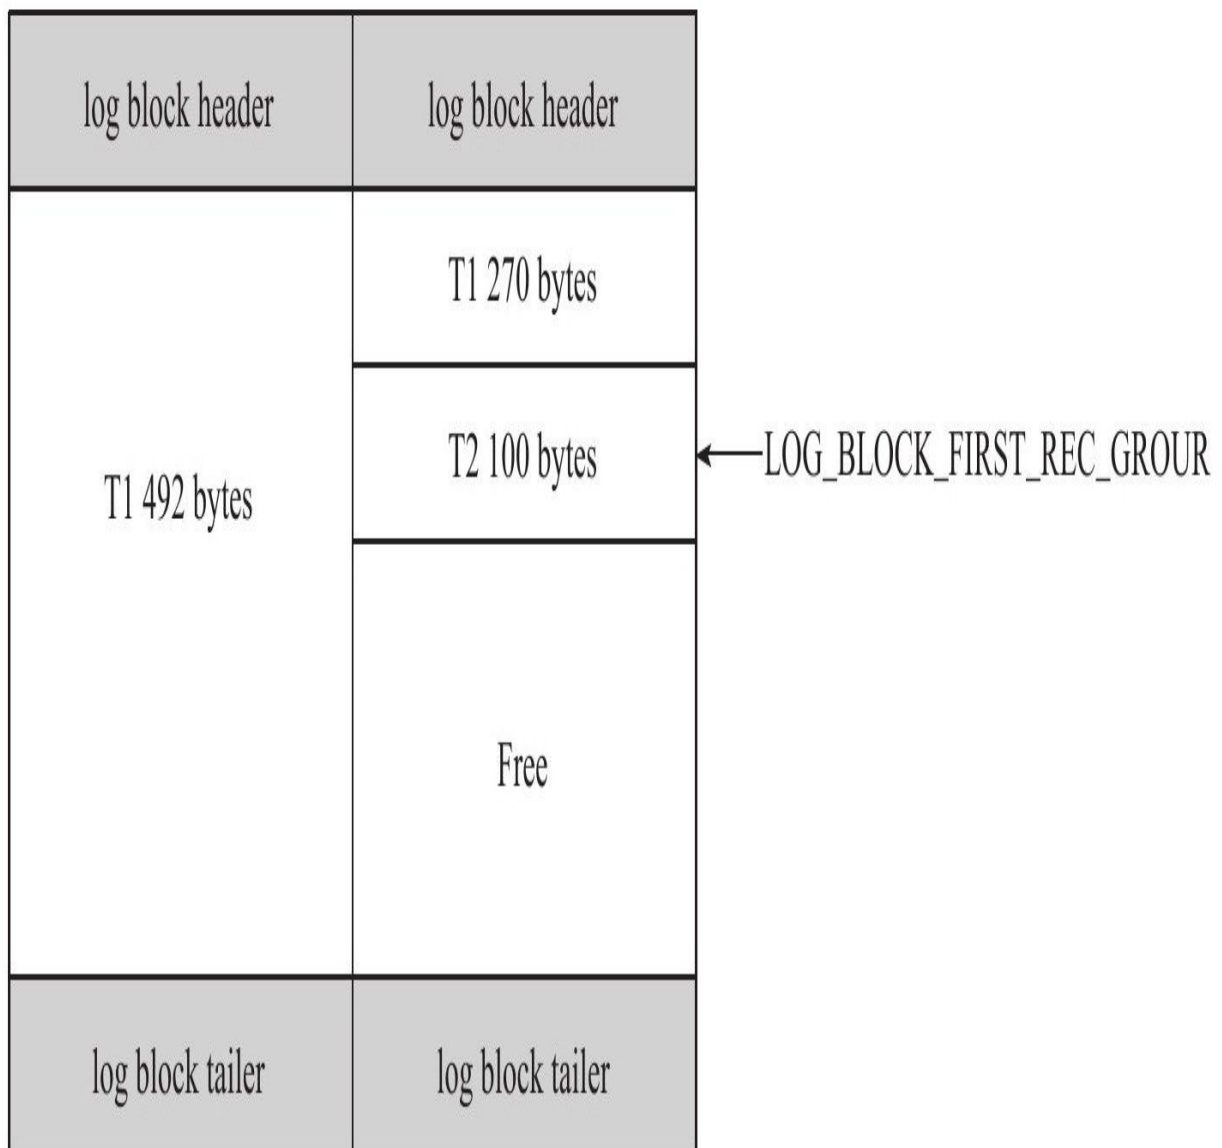
表 7-2 log block header

| 名 称 | 占用字节 |
|---------------------------|------|
| LOG_BLOCK_HDR_NO | 4 |
| LOG_BLOCK_HDR_DATA_LEN | 2 |
| LOG_BLOCK_FIRST_REC_GROUP | 2 |
| LOG_BLOCK_CHECKPOINT_NO | 4 |

```
log buffer[]log block[][][]log buffer[][][]
LOG_BLOCK_HDR_NO[][][]4[]
[]flush bit[][]2G[]
```

```
LOG_BLOCK_HDR_DATA_LEN 2 log block log
block 0x200 log block 512
```

```
LOG_BLOCK_FIRST_REC_GROUP[2]log block
LOG_BLOCK_HDR_DATA_LENlog
blockT11762T2100
log block492log buffer7-8
```



7-8 LOG_BLOCK_FIRST_REC_GROUP

7-8 T1 792 log block
 log block LOG_BLOCK_FIRST_REC_GROUP 12 log block
 log block T1 T2 log
 block LOG_BLOCK_FIRST_REC_GROUP
 282 270+12

LOG_BLOCK_CHECKPOINT_NO 4 log block
 4

log block tailer 1 7-3
LOG_BLOCK_HDR_NO log_block_init

表 7-3 log block tailer 部分

| 名 称 | 大小 (字节) |
|------------------|---------|
| LOG_BLOCK_TRL_NO | 4 |

3.log group

log group ha_innobase.cc InnoDB log group

log group log group log group InnoDB 1.2 4GB 4GB InnoDB 1.2 512GB InnoDB 1.1 4GB

log buffer log block log block 512 InnoDB log buffer log block

log buffer

log checkpoint

log block append redo log file redo log file redo log file round-robin

```

log blockredo log fileredo log file
redo log filelog bufferlog
block2KBredo log file2KB
log blocklog groupredo log file2KB
45127-4

```

表 7-4 redo log file 前 2KB 部分的内容

| 名 称 | 大小 (字节) |
|-----------------|---------|
| log file header | 512 |
| checkpoint1 | 512 |
| 空 | 512 |
| checkpoint2 | 512 |

```

log groupredo log filelog
groupredo log file
redo log filelog block2KB
InnoDBlog groupredo
log file7-9

```


Log Group 1

Redo Log File1

[illegible]

Redo Log File2

[illegible]

Log Group 2

Redo Log File1

[illegible]

Redo Log File2

[illegible]

7-9 log group redo log file

log file header InnoDB checkpoint
checkpoint

4.

InnoDB
7-10

| redo_log_type | space | page_no | redo log body |
|---------------|-------|---------|---------------|
|---------------|-------|---------|---------------|

7-10

3

redo_log_type

space ID

page_no

redo log body
7-11

MLOG_REC_INSERT

| | | | | | | | | |
|------|-------|---------|--------------------|---------------------|-----------|-------------------|---------------------|----------|
| type | space | page_no | cur rec
_offset | len &
extra_info | into_bits | origin_
offset | mis_mate
h_index | rec body |
|------|-------|---------|--------------------|---------------------|-----------|-------------------|---------------------|----------|

MLOG_REC_DELETE

| | | | |
|------|-------|---------|--------|
| type | space | page_no | offset |
|------|-------|---------|--------|

7-11

InnoDB1.251

5.LSN

LSN Log Sequence Number InnoDB
LSN 8 LSN

□ □ □ □ □ □ □ □ □ □

☐ checkpoint

5/11/2016

```
LSN[00000000000000000000000000000000]LSN[1 0000000000000000]T1[000
10000000000000000000000000000000]LSN[0000110000000000]T2[0002000000000000]
LSN[00001 300000]LSN[00000000000000000000000000000000]
```

LSN
FIL_PAGE_LSN
LSN
LSN
P1
LSN
10 000
InnoDB
LSN
13 000
P1
LSN
P1
LSN
P1
LSN

SHOW ENGINE INNODB STATUS

```
mysql>SHOW ENGINE INNODB STATUS\G;
```

```
-----
---
LOG
---
Log sequence number 11 3047174608

Log flushed up to 11 3047174608

Last checkpoint at 11 3047174608

0 pending log writes,0 pending chkp writes

142 log i/o's done,0.00 log i/o's/second

-----

1 row in set(0.00 sec)
```

Log sequence number
LSN
Log flushed up to
LSN
Last checkpoint at
LSN

Log sequence number
Log flushed up to
Last checkpoint at

```
mysql>show engine innodb status\G;
```

```
---
```

```
LOG
```

Log sequence number 203318213447

Log flushed up to 203318213326

Last checkpoint at 203252831194

1 pending log writes,0 pending chkp writes

103447 log i/o's done,7.00 log i/o's/second

.....

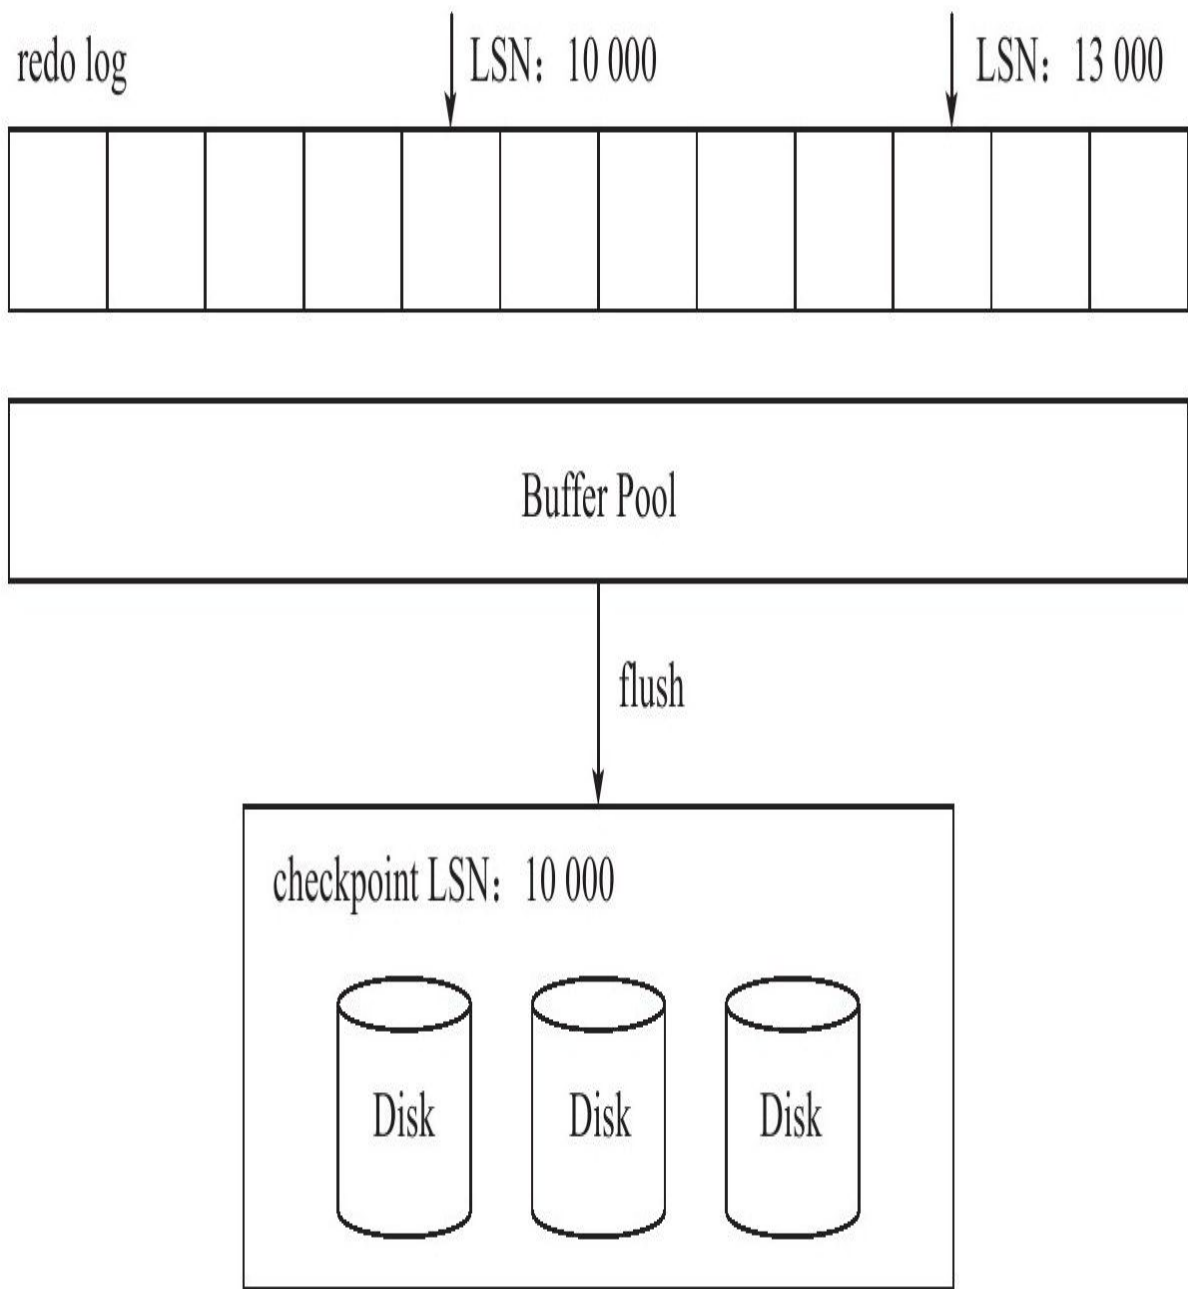
1 row in set(0.00 sec)

Log sequence number Log flushed up to Last
checkpoint at

6.

InnoDB
InnoDB
InnoDB

checkpoint LSN checkpoint
7-12 checkpoint LSN 10 000
LSN 10 000 13 000



7-12

InnoDB INSERT

```
CREATE TABLE t(a INT,b INT,PRIMARY KEY(a),KEY(b));
```

SQL

INSERT INTO t SELECT 1,2;

page(2,3),offset 32,value 1,2#

page(2,4),offset 64,value 2#

B+split

$$f(f(x))=f(x)$$

DBA ROW
INSERT
INSERT

7.2.2 undo

1. 简介

MySQL 数据库在 InnoDB 引擎下，redo 和 undo 是数据库的两个重要组件。redo 用于记录事务的修改，而 undo 用于回滚事务。当执行 ROLLBACK 操作时，数据库会使用 undo 日志来恢复数据到事务开始前的状态。

redo 和 undo 是数据库的两个重要组件。redo 用于记录事务的修改，而 undo 用于回滚事务。当执行 ROLLBACK 操作时，数据库会使用 undo 日志来恢复数据到事务开始前的状态。

```
[root@xen-server ~]#python py_innodb_page_info.py/usr/local/mysql/data/ibdata1
```

Total number of page:46208:

Insert Buffer Free List:13093

Insert Buffer Bitmap:3

System Page:5

Transaction system Page:1

Freshly Allocated Page:4579

undo Log Page:2222

File Segment inode:6

B-tree Node:26296

File Space Header:1

ibdata1:2

undo 日志是 InnoDB 引擎的一个重要组件，用于记录事务的修改。当执行 ROLLBACK 操作时，数据库会使用 undo 日志来恢复数据到事务开始前的状态。undo 日志的格式如下：

redo log
redo log

INSERT 10W
ROLLBACK
InnoDB
DELETE
DELETE
InnoDB
INSERT
UPDATE
InnoDB
UPDATE

undo
MVCC
InnoDB
MVCC
undo
undo

undo log
redo log
undo log
redo log
undo log

2.undo

InnoDB
undo
InnoDB
rollback segment
1024
undo log
segment
undo log segment
undo
5
0
5
rollback segment header
FIL_PAGE_TYPE_SYS

InnoDB1.1
1.1
rollback segment
1024
1.1
InnoDB
128
rollback segment
128*1024

InnoDB1.1
128
rollback segment
rollback segment
InnoDB1.2
rollback segment

innodb_undo_directory

innodb_undo_logs

innodb_undo_tablespaces

innodb_undo_directory rollback segment
rollback segment
“.”InnoDB

innodb_undo_logs rollback segment 128
InnoDB1.2 innodb_rollback_segments

innodb_undo_tablespaces rollback segment
rollback segment
innodb_undo_directory undo rollback
segment 7-13 3 rollback segment

```
mysql> SHOW VARIABLES LIKE 'innodb_undo%';
```

| Variable_name | Value |
|-------------------------|-------|
| innodb_undo_directory | . |
| innodb_undo_logs | 128 |
| innodb_undo_tablespaces | 3 |

```
3 rows in set (0.00 sec)
```

```
mysql> SHOW VARIABLES LIKE 'datadir';
```

| Variable_name | Value |
|---------------|-------------------------------|
| datadir | /Users/david/mysql_data/data/ |

```
1 row in set (0.00 sec)
```

```
mysql> system ls -lh/Users/david/mysql_data/data/undo*
```

```
-rw-rw---- 1 david staff 10M 11 22 16:55/Users/david/mysql_data/data/undo001
-rw-rw---- 1 david staff 10M 11 22 16:51/Users/david/mysql_data/data/undo002
-rw-rw---- 1 david staff 10M 11 22 16:51/Users/david/mysql_data/data/undo003
```

7-13 3 rollback segment

undo log segment undo log
InnoDB

undo log purge

undo log

undo log undo log
undo log undo log
undo log undo log purge

undo OLTP
TPS transaction per
second 1000 undo 1000*60
1GB purge 20
InnoDB undo undo log
undo 3/4 undo
undo log undo log
undo undo log purge
undo log

SHOW ENGINE INNODB STATUS undo log

```
mysql> SHOW ENGINE INNODB STATUS\G;
```

```
*****1.row*****
```

```
.....
```

```
-----
```

```
TRANSACTIONS
```

```
-----
```

```
Trx id counter 3000
```

```
Purge done for trx's n:o2C03 undo n:o0
```

```
History list length 12
```

LIST OF TRANSACTIONS FOR EACH SESSION:

---TRANSACTION 0,not started

MySQL thread id 1,OS thread handle 0x1500f1000,query id 4 localhost root

show engine innodb status

.....

History list length 12 purge
undo logHistory list length
0

3.undo log

InnoDBundo log

insert undo log

update undo log

insert undo loginsertundo loginsert
undo logundo log
purgeinsert undo log7-14

insert undo log record

n_unique_index

| | |
|-----------|-------|
| next | |
| type_cmpl | |
| *undo no | |
| *table id | |
| *len1 | col1 |
| *len2 | col2 |
| | |
| *lenN | colN |
| start | |

7-14 insert undo log

7-14 insert undo log
*
insert undo log
next
undo log
next
undo log
type_cmpl
undo
insert undo log
11
undo_no
ID
table_id
undo log
rollback
undo log

update undo log
delete
update
undo log
undo log
MVCC
undo log
purge
update undo log
7-15

| | | | |
|----------------|------------------------|--------|------------|
| | update undo log record | | |
| | next | | |
| | type_cmpl | | |
| | *undo no | | |
| | *table id | | |
| | info_bits | | |
| | *DATA_TRX_ID | | |
| | *DATA_ROLL_PTR | | |
| n_unique_index | *len1 | i_col1 | |
| | *len2 | i_col2 | |
| | | | |
| | *lenN | i_colN | |
| update vector | n_update_field | | |
| | *pos1 | *len1 | u_old_col1 |
| | *pos2 | *len2 | u_old_col2 |
| | | | |
| | *posN | *lenN | u_old_colN |
| | n_bytes_below | | |
| | *pos | *len | col1 |
| | *pos | *len | col2 |
| | | | |
| | *pos | *len | colN |
| | start | | |

7-15 update undo log

update undo log insert undo log
next start undo_no table_id insert undo log
type_cmpl update undo log

12 TRX_UNDO_UPD_EXIST_REC non-delete-mark

13 TRX_UNDO_UPD_DEL_REC delete not delete

14 TRX_UNDO_DEL_MARK_REC delete

update_vector update_vector update
undo log undo log
undo

4. undo

Oracle Microsoft SQL Server undo
InnoDB DBA InnoDB
information_schema undo

INNODB_TRX_ROLLBACK_SEGMENT
rollback segment 7-16

```
mysql> DESC INNODB_TRX_ROLLBACK_SEGMENT;
```

| Field | Type | Null | Key | Default | Extra |
|--------------------|---------------------|------|-----|---------|-------|
| Segment_id | bigint(21) unsigned | NO | | 0 | |
| space | bigint(21) unsigned | NO | | 0 | |
| page_no | bigint(21) unsigned | NO | | 0 | |
| last_page_no | bigint(21) unsigned | YES | | NULL | |
| last_offset | bigint(21) unsigned | NO | | 0 | |
| last_trx_no | varchar(18) | NO | | | |
| update_undo_list | bigint(21) unsigned | NO | | 0 | |
| update_undo_cached | bigint(21) unsigned | NO | | 0 | |
| insert_undo_list | bigint(21) unsigned | NO | | 0 | |
| insert_undo_cached | bigint(21) unsigned | NO | | 0 | |

```
10 rows in set (0.00 sec)
```

□ 7-16 INNODB_TRX_ROLLBACK_SEGMENT □□□

查看innodb rollback segment

```
mysql>SELECT segment_id,space,page_no
```

```
->FROM INNODB_TRX_ROLLBACK_SEGMENT;
```

```
+-----+-----+-----+
```

```
|segment_id|space|page_no|
```

```
+-----+-----+-----+
```

```
|0|0|6|
```

```
|1|0|45|
```

```
|2|0|46|
```

```
.....
```

```
128 rows in set(0.00 sec)
```

查看innodb INNODB_TRX_UNDO 空间大小及undo log 大小
innodb 空间大小
innodb INNODB_TRX_UNDO 空间大小
t

```
CREATE TABLE t(
```

```
a INT,
```

```
b VARCHAR(32),
```

```
PRIMARY KEY(a),
```

```
KEY(b)
```

```
)ENGINE=InnoDB;
```

查看innodb INNODB_TRX_UNDO 空间大小及undo log 大小

```
mysql>BEGIN;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>INSERT INTO t SELECT 1,'1';
```

```
Query OK,1 row affected(0.00 sec)
```

Records:1 Duplicates:0 Warnings:0

mysql>SELECT*FROM information_schema.INNODB_TRX_UNDO\G;

*****1.row*****

trx_id:3001

rseg_id:2

undo_rec_no:0

undo_rec_type:TRX_UNDO_INSERT_REC

size:12

space:0

page_no:334

offset:272

1 row in set(0.00 sec)

undo log ID 3001 rollback segment ID 2
undo_rec_no 0
TRX_UNDO_INSERT_REC insert undo log size undo log
12 space page_no offset undo log
ibdata1 334 272 12

01 1c 0b 00 16 04 80 00 00 01 01 10

undo log undo log

01 1c#undo log 272+12=0x011c

0b#undo log TRX_UNDO_INSERT_REC 11

00#undo log undo_rec_no

16# ID

04#

80 00 00 01#

01 10#undo log 272=0x0110

mysql>select undo_log,rollback segment_id=2 from information_schema.INNODB_TRX_ROLLBACK_SEGMENT where rollback segment=1

```
mysql>SELECT segment_id,insert_undo_list,insert_undo_cached
```

```
-FROM information_schema.INNODB_TRX_ROLLBACK_SEGMENT
```

```
-WHERE segment_id=2\G;
```

```
*****1.row*****
```

```
segment_id:2
```

```
insert_undo_list:1
```

```
insert_undo_cached:0
```

```
1 row in set(0.00 sec)
```

mysql>insert_undo_list=1 commit

```
mysql>COMMIT;
```

```
Query OK, 0 rows affected(0.00 sec)
```

```
mysql>SELECT segment_id,insert_undo_list,insert_undo_cached
```

```
-FROM information_schema.INNODB_TRX_ROLLBACK_SEGMENT
```

```
-WHERE segment_id=2\G;
```

```
*****1.row*****
```

```
segment_id:2
```

```
insert_undo_list:0
```

```
insert_undo_cached:1
```

```
1 row in set(0.00 sec)
```

mysql>insert_undo_list=0 insert_undo_cached=1
undo,rollback segment undo

mysql>delete undo_log

mysqlBEGIN;

Query OK,0 rows affected(0.00 sec)

mysqlDELETE FROM t WHERE a=1;

Query OK,1 row affected(0.00 sec)

Records:1 Duplicates:0 Warnings:0

mysqlSELECT*FROM information_schema.INNODB_TRX_UNDO\G;

*****1.row*****

trx_id:3201

rseg_id:2

undo_rec_no:0

undo_rec_type:TRX_UNDO_DEL_MARK_REC

size:37

space:0

page_no:326

offset:620

1 row in set(0.00 sec)

0000000000000000000032600000006200000000000000000

0518260 00 00 00 00 00 00 00 00 00 00 00 02 91 0e 00

0518270 16 00 00 00 00 30 01 e0 82 00 00 01 4e 01 10 04

0518280 80 00 00 01 00 0b 00 04 80 00 00 01 03 01 31 02

0518290 6c 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

00000000

02 91#undo log

0e#undo logTRX_UNDO_DEL_MARK_REC14

00#undo no

```
16#table id

00#info bits

00 00 00 30 01 e0#recid

82 00 00 01 4e 01 10#rec

04#

80 00 00 01#

00 0b#

00#

04#

80 00 00 01#

03#00000000002

01#

31#b'1'

02 6c#
```

rollback segment

```
mysql>SELECT segment_id,update_undo_list,update_undo_cached

- FROM information_schema.INNODB_TRX_ROLLBACK_SEGMENT

- WHERE segment_id=2\G;

*****1.row*****

segment_id:2

update_undo_list:1

update_undo_cached:0

1 row in set(0.00 sec)
```

undo cache

```
mysql>COMMIT;

Query OK,0 rows affected(0.00 sec)
```

```
mysql>SELECT segment_id,update_undo_list,update_undo_cached
```

```
->FROM information_schema.INNODB_TRX_ROLLBACK_SEGMENT
```

```
->WHERE segment_id=2\G;
```

```
*****1.row*****
```

```
segment_id:2
```

```
update_undo_list:0
```

```
update_undo_cached:1
```

```
1 row in set(0.00 sec)
```

delete flag 1 purge

update undo log 1 '1' update
INNODB_TRX_UNDO undo log

```
mysql>INSERT INTO t SELECT 1,'1';
```

```
mysql>BEGIN;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>UPDATE t SET b='2'WHERE a=1;
```

```
Query OK,1 row affected(0.00 sec)
```

```
Rows matched:1 Changed:1 Warnings:0
```

```
mysql>SELECT*FROM information_schema.INNODB_TRX_UNDO\G;
```

```
*****1.row*****
```

```
trx_id:3205
```

```
rseg_id:5
```

```
undo_rec_no:0
```

```
undo_rec_type:TRX_UNDO_UPD_EXIST_REC
```

```
size:41
```

```
space:0
```

```
page_no:318
```

```
offset:724
```


1 row in set(0.00 sec)

318724

04f82d0 00 00 00 02 fd 0c 00 16 00 00 00 00 32 04 e0

04f82e0 84 00 00 01 48 01 10 04 80 00 00 01 01 03 01 31

04f82f0 00 0b 00 04 80 00 00 01 03 01 31 02 d4 00 00 00

02 fd#undo log

0c#undo logTRX_UNDO_UPD_DEL_REC13

00#undo no

16#table id

00#info bits

00 00 00 32 04 e0#rec trx id

84 00 00 01 48 01 10#rec

04#

80 00 00 01#

01#update vector

03#update vectorb

01#update vector

31#update vector'1'

00 0b#

00#

04#

80 00 00 01#

03#

31#

02 d4#undo log

undo log

mysql>ROLLBACK;

Query OK, 1 row affected (0.00 sec)

mysql>UPDATE t SET a=2 WHERE a=1;

Rows matched: 1 Changed: 1 Warnings: 0

mysql>SELECT*FROM information_schema.INNODB_TRX_UNDO

->ORDER BY undo_rec_no\G;

*****1.row*****

trx_id:320F

rseg_id:11

undo_rec_no:0

undo_rec_type:TRX_UNDO_DEL_MARK_REC

size:37

space:0

page_no:324

offset:492

*****2.row*****

trx_id:320F

rseg_id:11

undo_rec_no:1

undo_rec_type:TRX_UNDO_INSERT_REC

size:12

space:0

page_no:336

offset:272

2 rows in set (0.00 sec)

update
TRX_UNDO_DEL_MARK_REC undo log
TRX_UNDO_INSERT_REC undo log undo_rec_no
undo log

InnoDB undo DBA
undo

7.2.3 purge

delete/update操作は、MySQLのSQL文で

```
DELETE FROM t WHERE a=1;
```

tのa列の値がbのdelete操作のundo logは、delete flagが1のB+ツリーでa=1、b=1のundo logは、purge操作で

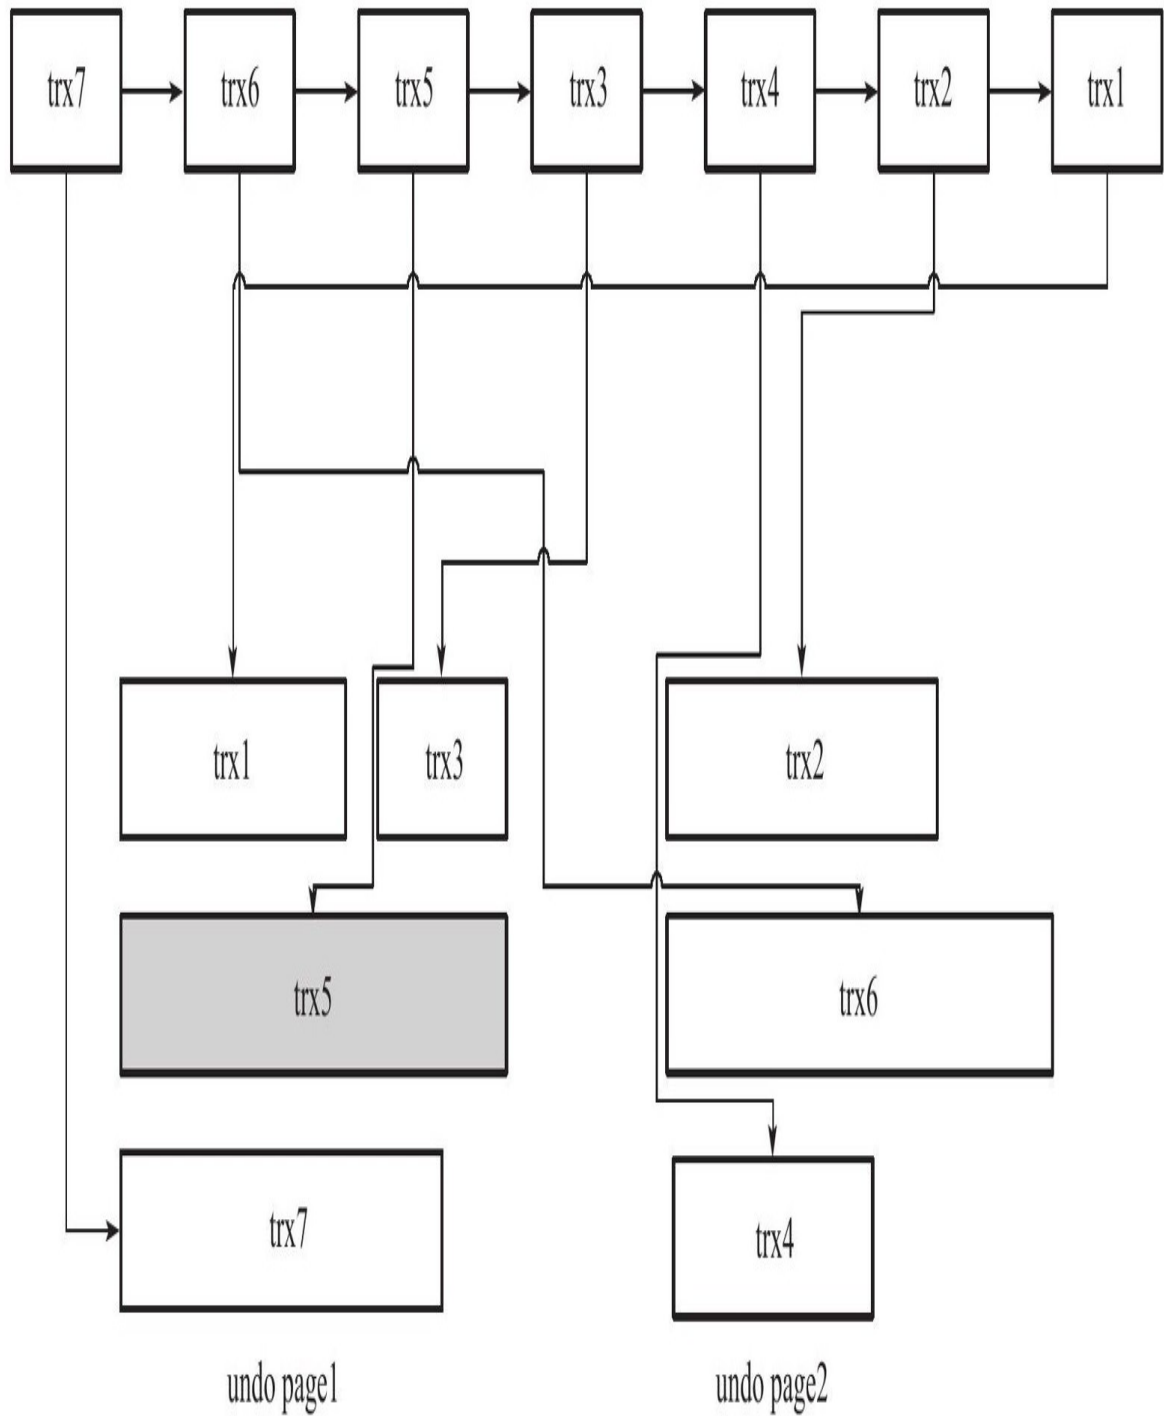
purge操作はdelete/update操作のInnoDBのMVCC操作のInnoDBのpurge操作はdelete/update操作の“”delete操作で

InnoDBのundo logは、undo logは、InnoDBのhistory undo logは、

7-17 history list undo log InnoDB undo page undo log undo page undo log trx5 undo log

purge InnoDB history list trx1 InnoDB trx1 undo log trx3 trx5 trx5 history list trx2 trx2 trx6 trx4 undo page2 undo page

History List



7-17 undo log history

InnoDB history list undo log undo page
undo log purge

innodb_purge_batch_size purge undo
page InnoDB1.2 20 1.2
300 undo page undo
page purge
undo page CPU IO undo log
DBA MySQL

InnoDB purge history list
innodb_max_purge_lag history list
“ ” DML 0 history list
0 DML

```
delay=length(history_list)-innodb_max_purge_lag*10)-5
```

delay delay DML
update 5 delay 5*delay
delay purge

InnoDB1.2 innodb_max_purge_lag_delay
delay delay delay
innodb_max_purge_lag_delay purge SQL

7.2.4 group commit

```

000000000000000000000000fsync000000000000000000000000
0000000000000000000000000000000000000000000fsync0000
00000000fsync000000000000group commit000000fsync00
000000000000000000InnoDB000000000000000000000000

```

1

2fsync

[illegible]

```

InnoDB1.2
InnoDBgroup commit
replication

```

[illegible]

1 InnoDB prepare

2 MySQL

3 InnoDB

[illegible]

bfsync

[illegible]

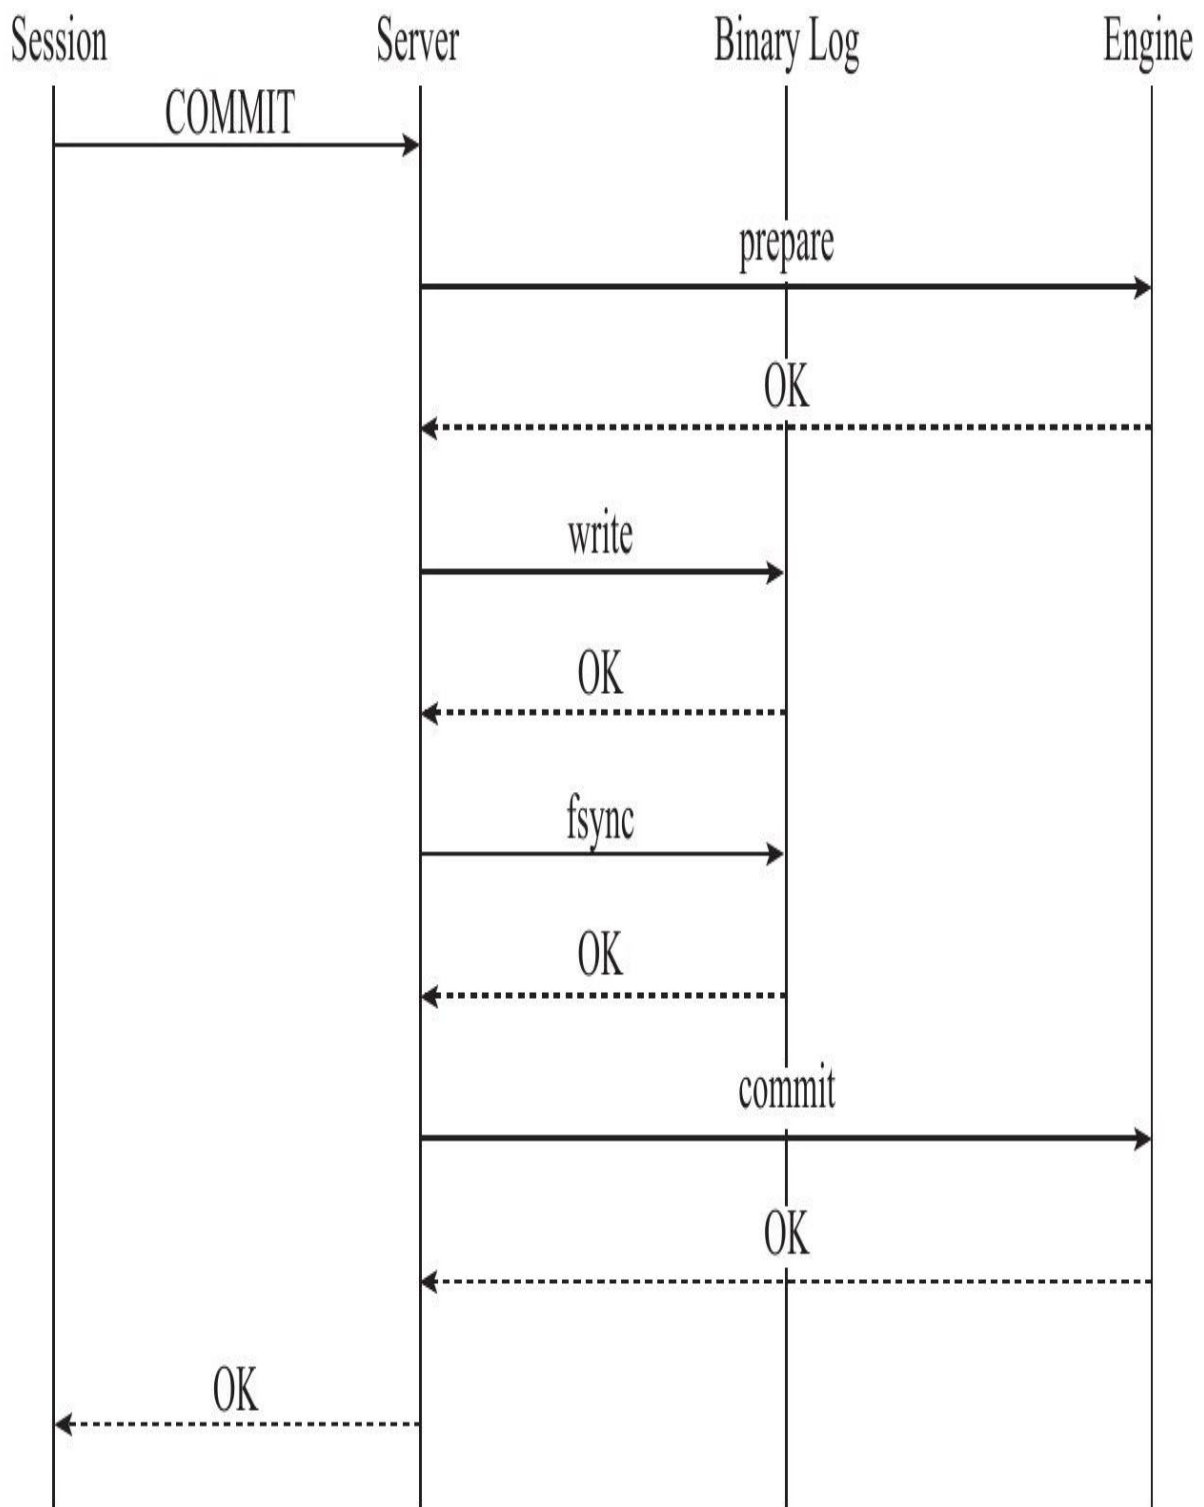
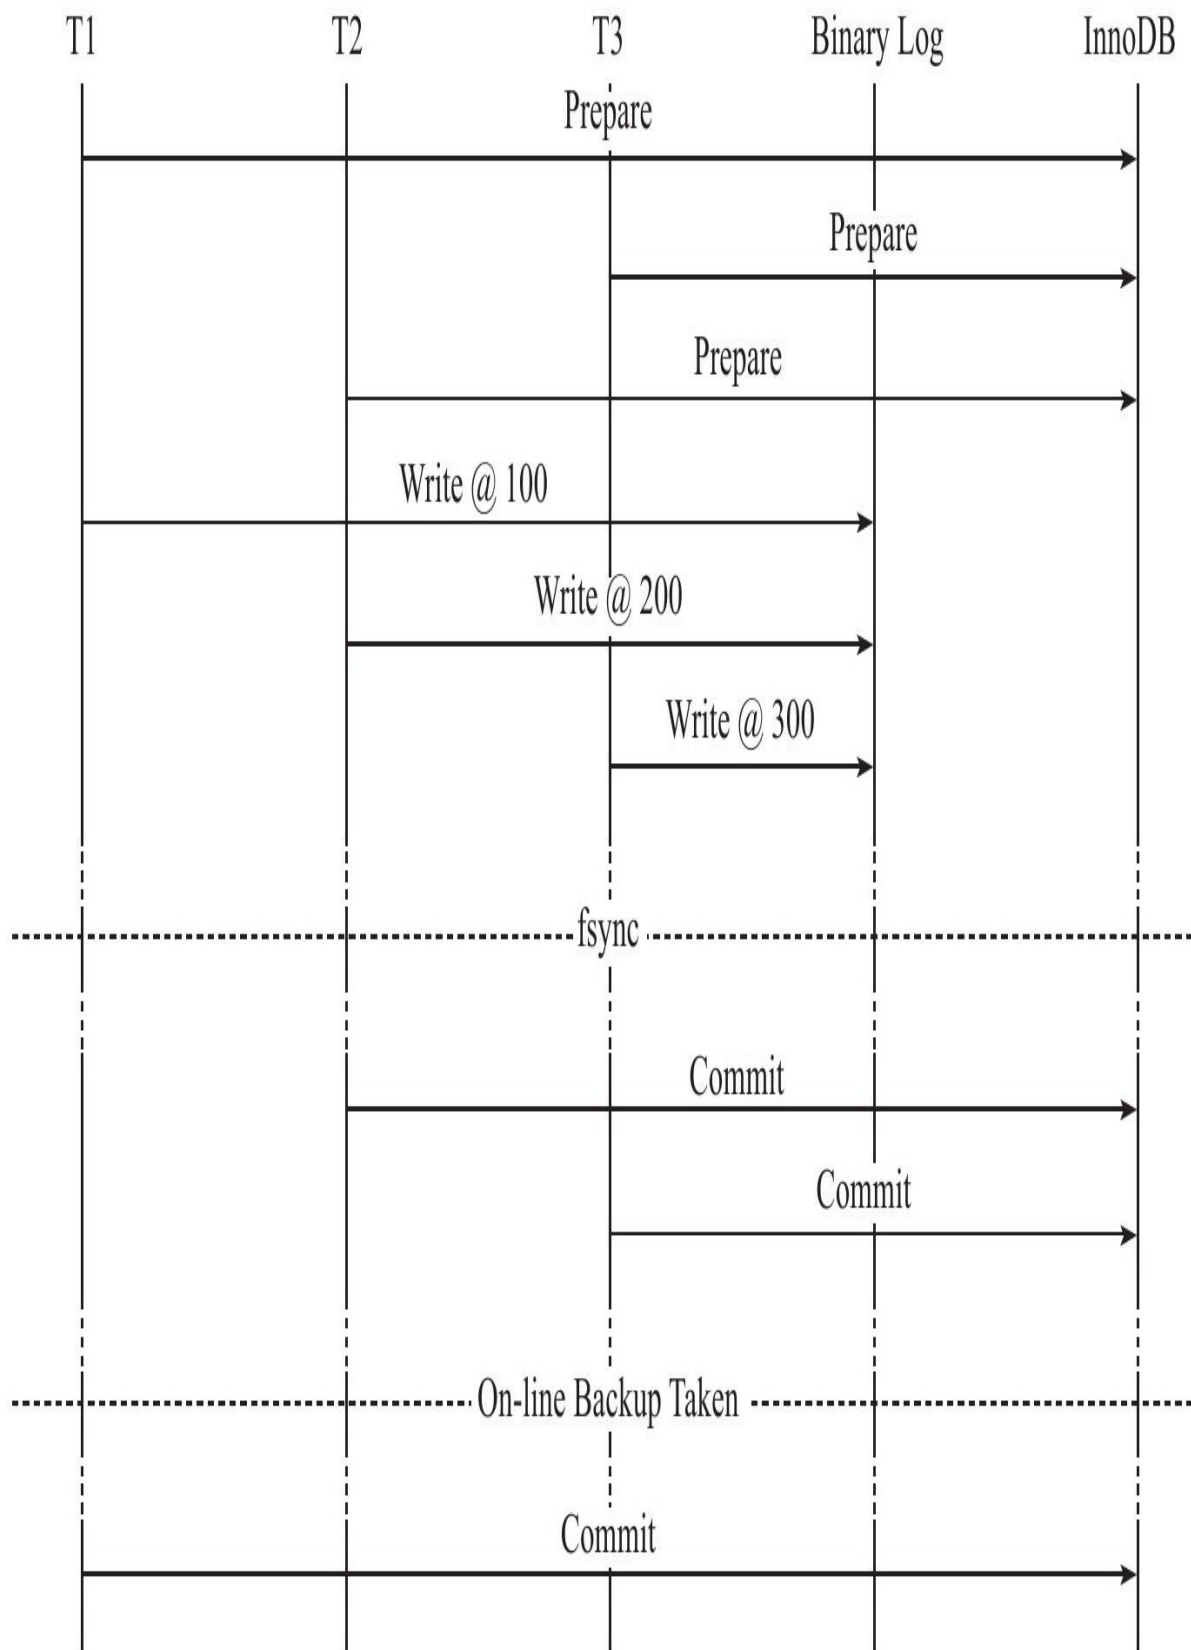
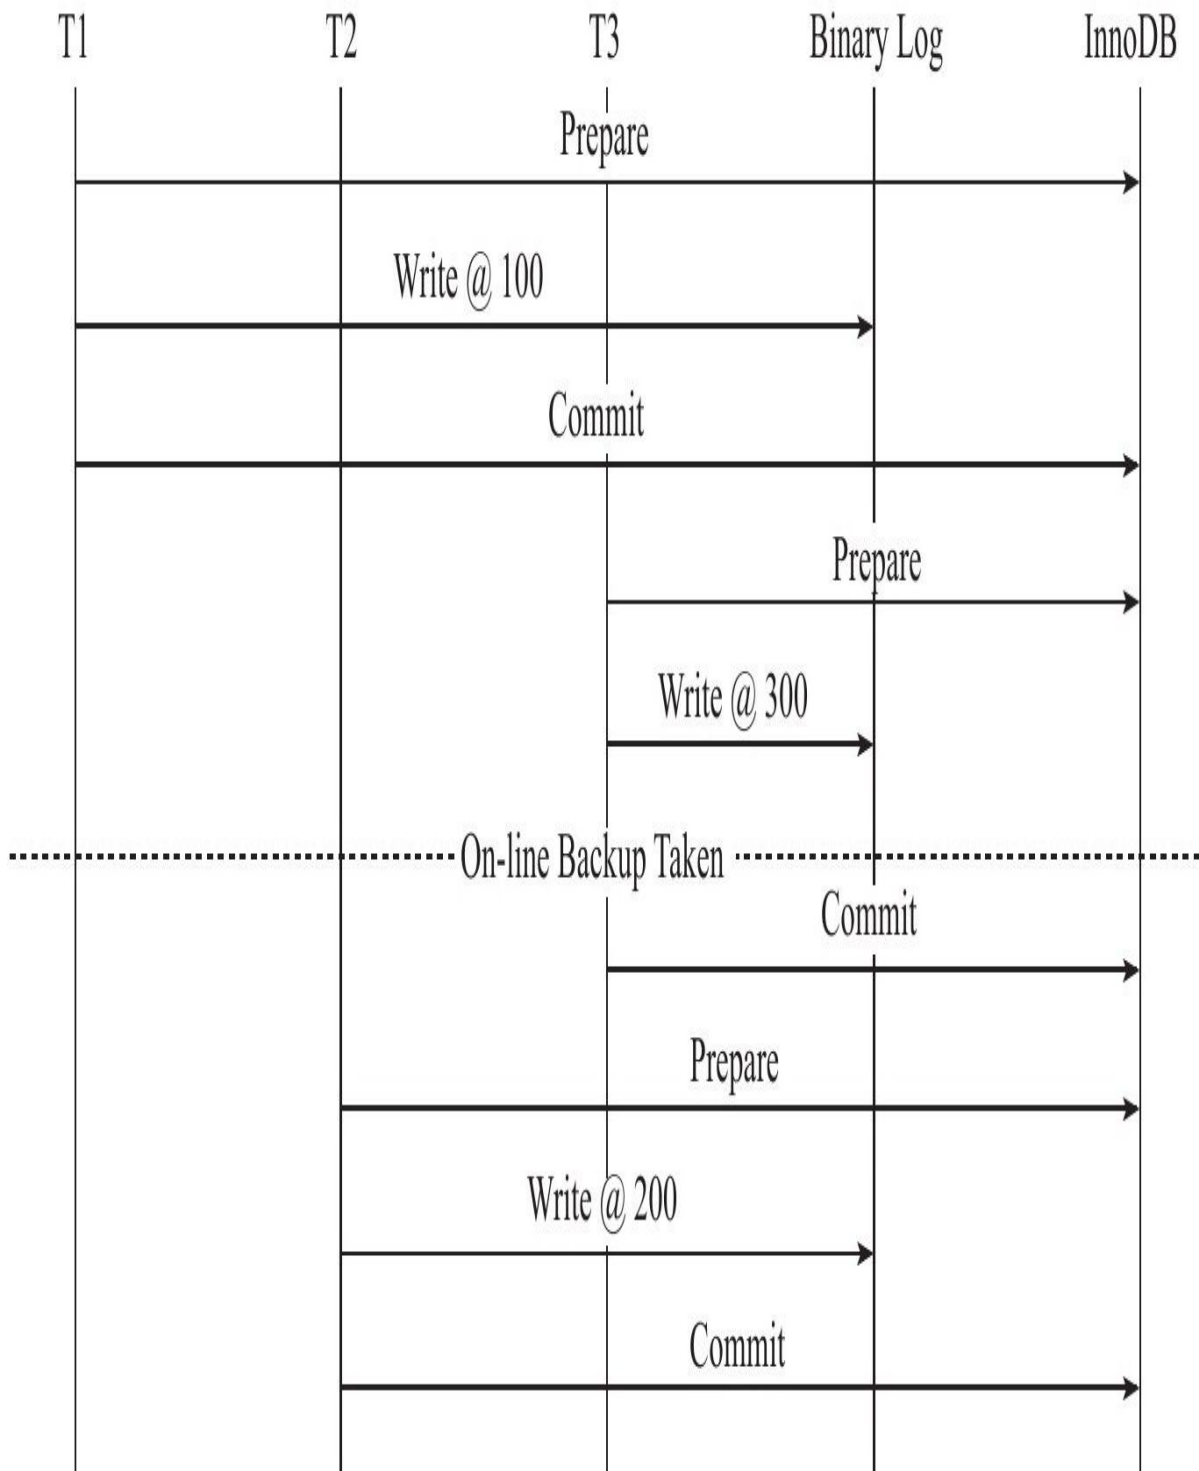


图 7-18 InnoDB 的 commit 过程

MySQL InnoDB MySQL
prepare_commit_mutex 3 a
b group commit

MySQL InnoDB
xtrabackup ibbackup
replication 7-19

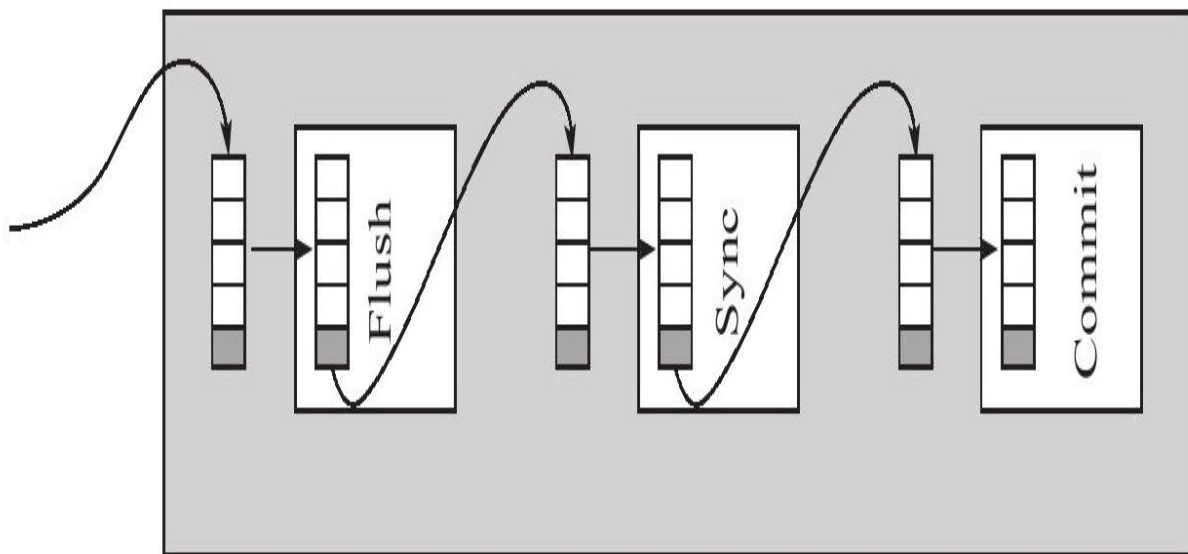




7-20 prepare_commit_mutex InnoDB MySQL

2010 MySQL Facebook MySQL Percona MariaDB Kristian Nielsen “” MySQL group commit InnoDB group commit prepare_commit_mutex MySQL 5.6 Binary Log Group Commit BLGC

MySQL 5.6 BLGC 7-21



7-21 MySQL 5.6 BLGC

MySQL leader follower leader follower BLGC

Flush

Sync fsync BLGC

Commit leader InnoDB
group commit prepare_commit_mutex group
commit

Commit Flush group commit
group commit
group commit

binlog_max_flush_queue_time Flush
Sync
group commit 0
0 MySQL 100

7.3 数据库事务

MySQL数据库默认是开启auto commit的，即SQL语句执行后，会自动提交。如果希望在一个事务中执行多条SQL语句，则需要使用BEGIN或START TRANSACTION来开启事务，使用COMMIT来提交事务，使用ROLLBACK来回滚事务。在MySQL中，可以通过SET AUTOCOMMIT=0来关闭自动提交功能。

START TRANSACTION|BEGIN 开启事务

COMMIT 提交事务
COMMIT WORK 提交事务
COMMIT 提交事务

ROLLBACK 回滚事务
ROLLBACK WORK 回滚事务
回滚

SAVEPOINT identifier:SAVEPOINT 设置保存点
SAVEPOINT

RELEASE SAVEPOINT identifier 释放保存点
释放保存点

ROLLBACK TO[SAVEPOINT]identifier 回滚到保存点
回滚到保存点
UPDATE
SAVEPOINT
DELETE
DELETE
ROLLBACK TO SAVEPOINT
SAVEPOINT
DELETE
UPDATE

SET TRANSACTION 设置事务隔离级别
InnoDB
READ UNCOMMITTED
READ COMMITTED
REPEATABLE
READ
SERIALIZABLE

START TRANSACTION
BEGIN
MySQL
MySQL
BEGIN...END
START TRANSACTION

COMMIT
COMMIT WORK
COMMIT WORK
CHAIN
RELEASE
CHAIN

completion_type
0
COMMIT
COMMIT WORK
completion_type
1
COMMIT WORK
COMMIT AND CHAIN

```
mysql>CREATE TABLE t(a INT,PRIMARY KEY(a))ENGINE=INNODB;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>SELECT @@autocommit\G;
```

```
*****1.row*****
```

```
@@autocommit:1
```

```
1 row in set(0.00 sec)
```

```
mysql>SET @@completion_type=1;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>BEGIN;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>INSERT INTO t SELECT 1;
```

```
Query OK,1 row affected(0.00 sec)
```

```
Records:1 Duplicates:0 Warnings:0
```

```
mysql>COMMIT WORK;
```

```
Query OK,0 rows affected(0.01 sec)
```

```
mysql>INSERT INTO t SELECT 2;
```

```
Query OK,1 row affected(0.00 sec)
```

```
Records:1 Duplicates:0 Warnings:0
```

```
mysql>INSERT INTO t SELECT 2;
```


ERROR 1062(23000):Duplicate entry'2'for key'PRIMARY'

mysql>ROLLBACK;

Query OK,0 rows affected(0.00 sec)

#00000000100000000020000

mysql>SELECT*FROM t\G;

*****1.r0w*****

a:1

1 row in set(0.00 sec)

0000000000completion_type01000000COMMIT WORK000100
000000000020000000BEGIN000START TRANSACTION00000000
0000000000000000002000000000ROLLBACK0000000001000
000200000000completion_type0100COMMIT WORK00000000
00000000INSERT INTO t SELECT 20000000000000000002000000
00000000t00

00completion_type0200COMMIT WORK000COMMIT AND
RELEASE000000000000000000000000

mysql>SET@@completion_type=2;

Query OK,0 rows affected(0.00 sec)

mysql>BEGIN;

Query OK,0 rows affected(0.00 sec)

mysql>INSERT INTO t SELECT 3;

Query OK,1 row affected(0.00 sec)

Records:1 Duplicates:0 Warnings:0

mysql>COMMIT WORK;

Query OK,0 rows affected(0.01 sec)

mysql>SELECT@@version\G;

ERROR 2006(HY000):MySQL server has gone away

No connection.Trying to reconnect...

Connection id:54

Current database:test

*****1.row*****

@@version:5.1.45-log

1 row in set(0.00 sec)

completion_type2COMMIT WORK
SELECT@@versionERROR 2006HY000MySQL
server has gone awayCOMMIT
WORK

ROLLBACKROLLBACK WORKCOMMITCOMMIT WORK

SAVEPOINTROLLBACK TO SAVEPOINT

mysqlBEGIN;

Query OK,0 rows affected(0.00 sec)

mysqlROLLBACK TO SAVEPOINT t1;

ERROR 1305(42000):SAVEPOINT t1 does not exist

InnoDB

mysqlCREATE TABLE t(a INT,PRIMARY KEY(a))ENGINE=INNODB;

Query OK,0 rows affected(0.00 sec)

mysqlBEGIN;

Query OK,0 rows affected(0.00 sec)

mysqlINSERT INTO t SELECT 1;

Query OK, 1 row affected (0.00 sec)

Records: 1 Duplicates: 0 Warnings: 0

```
mysql>INSERT INTO t SELECT 1;
```

ERROR 1062 (23000): Duplicate entry '1' for key 'PRIMARY'

```
mysql>SELECT *FROM t\G;
```

*****1. row*****

a: 1

1 row in set (0.00 sec)

1062错误，插入失败，需要回滚
COMMIT ROLLBACK

ROLLBACK TO SAVEPOINT ROLLBACK
ROLLBACK TO SAVEPOINT
COMMIT ROLLBACK

```
mysql>CREATE TABLE t(a INT,PRIMARY KEY(a))ENGINE=INNODB;
```

Query OK, 0 rows affected (0.00 sec)

```
mysql>BEGIN;
```

Query OK, 0 rows affected (0.00 sec)

```
mysql>INSERT INTO t SELECT 1;
```

Query OK, 1 row affected (0.00 sec)

Records: 1 Duplicates: 0 Warnings: 0

```
mysql>SAVEPOINT t1;
```

Query OK, 0 rows affected (0.00 sec)

```
mysql>INSERT INTO t SELECT 2;
```

Query OK, 1 row affected (0.00 sec)

Records: 1 Duplicates: 0 Warnings: 0

```
mysql>SAVEPOINT t2;
```

Query OK, 0 rows affected (0.00 sec)

```
mysql>RELEASE SAVEPOINT t1;
```

Query OK, 0 rows affected (0.00 sec)

```
mysql> INSERT INTO t SELECT 2;
```

ERROR 1062 (23000): Duplicate entry '2' for key 'PRIMARY'

```
mysql> ROLLBACK TO SAVEPOINT t2;
```

Query OK, 0 rows affected (0.00 sec)

```
mysql> SELECT * FROM t;
```

+---+

|a|

+---+

|1|

|2|

+---+

2 rows in set (0.00 sec)

```
mysql> ROLLBACK;
```

Query OK, 0 rows affected (0.00 sec)

```
mysql> SELECT * FROM t;
```

Empty set (0.00 sec)

mysql> ROLLBACK TO SAVEPOINT t2
mysql> ROLLBACK
mysql> ROLLBACK TO SAVEPOINT

7.4 数据库SQL

数据库SQL语句包括DDL、DML、DCL、TCL、COMMIT等。

DDL包括ALTER DATABASE...UPGRADE DATA DIRECTORY NAME、ALTER EVENT、ALTER PROCEDURE、ALTER TABLE、ALTER VIEW、CREATE DATABASE、CREATE EVENT、CREATE INDEX、CREATE PROCEDURE、CREATE TABLE、CREATE TRIGGER、CREATE VIEW、DROP DATABASE、DROP EVENT、DROP INDEX、DROP PROCEDURE、DROP TABLE、DROP TRIGGER、DROP VIEW、RENAME TABLE、TRUNCATE TABLE。

MySQL包括CREATE USER、DROP USER、GRANT、RENAME USER、REVOKE、SET PASSWORD。

包括ANALYZE TABLE、CACHE INDEX、CHECK TABLE、LOAD INDEX INTO CACHE、OPTIMIZE TABLE、REPAIR TABLE。

Microsoft SQL Server包括DDL、DML、DCL、TCL、COMMIT等。
Microsoft SQL Server包括DDL、DML、DCL、TCL、COMMIT等。
InnoDB、Oracle。

TRUNCATE TABLE、DDL、DELETE、Microsoft SQL Server。

```
mysql>SELECT*FROM t\G;
```

```
*****1.row*****
```

```
a:1
```

```
*****2.row*****
```

```
a:2
```

```
2 rows in set(0.00 sec)
```

```
mysql>BEGIN;
```

```
Query OK,0 rows affected(0.01 sec)
```

```
mysql>TRUNCATE TABLE t;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>ROLLBACK;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>SELECT*FROM t;
```

```
Empty set(0.00 sec)
```

7.5

InnoDB InnoDB
 Question Per Second QPS
 Transaction Per Second TPS

```
TPS[com_commit+com_rollback]/time
autocommit=1
com_commitcom_rollback
```

```
mysql> SHOW GLOBAL STATUS LIKE 'com_commit'\G;
```

```
*****1.row*****
```

Variable_name:Com_commit

Value:5

```
1 row in set(0.00 sec)
```

```
mysql>INSERT INTO t SELECT 3;
```

Query OK, 1 row affected (0.00 sec)

Records:1 Duplicates:0 Warnings:0

```
mysql> SELECT * FROM t\G;
```

```
*****1.row*****
```

a:1

```
*****2.row*****
```

a:2

```
*****3.row*****
```

a:3

```
3 rows in set (0.00 sec)
```

```
mysql> SHOW GLOBAL STATUS LIKE 'com_commit'\G;
```

```
*****1.row*****
```

Variable name:Com commit

Value:5

1 row in set (0.00 sec)

MySQL handler_commit handler_rollback
MySQL 5.1 InnoDB
InnoDB Plugin “”
com_commit
com_rollback

7.6 数据库事务

数据库事务是指一个数据库操作序列，这些操作要么都执行，要么都不执行，是一个不可分割的工作单位。ISO/ANSI SQL定义了四种事务隔离级别：READ UNCOMMITTED、REPEATABLE READ、SERIALIZABLE 和Serializable。

SQL事务隔离级别

□ READ UNCOMMITTED

□ READ COMMITTED

□ REPEATABLE READ

□ SERIALIZABLE

READ UNCOMMITTED 允许脏读、不可重复读和幻读。
READ COMMITTED 允许脏读和不可重复读，但不允许幻读。
REPEATABLE READ 允许脏读，但不允许不可重复读和幻读。
SERIALIZABLE 不允许脏读、不可重复读和幻读。
SQL2 定义了四种事务隔离级别：READ UNCOMMITTED、READ COMMITTED、REPEATABLE READ 和 SERIALIZABLE。

InnoDB 默认的事务隔离级别是 REPEATABLE READ。InnoDB 通过 Next-Key Lock 实现 REPEATABLE READ。Microsoft SQL Server 默认的事务隔离级别是 READ COMMITTED。InnoDB 的事务隔离级别可以通过 SQL 语句修改为 SERIALIZABLE。

事务隔离级别越高，数据的一致性越好，但性能也会越低。READ COMMITTED 是大多数数据库的默认隔离级别。

SERIALIZABLE 是最高级别的隔离，可以保证事务的串行化。Jim Gray 的《Transaction Processing: Concepts and Techniques》一书详细讨论了事务处理的概念和技术。SERIALIZABLE 的实现非常复杂，InnoDB 默认使用 REPEATABLE READ，但可以通过 SQL 语句修改为 READ COMMITTED 或 SERIALIZABLE。

❑ InnoDB ❶❷❸❹❺❻❼❽❾❿⓫⓬⓭⓮⓯⓰⓱⓲⓳⓴⓵⓶⓷⓸⓹⓺⓻⓼⓽⓾⓿

```
SET[GLOBAL|SESSION]TRANSACTION ISOLATION LEVEL
```

```
{  
  
  READ UNCOMMITTED  
  
  |READ COMMITTED  
  
  |REPEATABLE READ  
  
  |SERIALIZABLE  
  
}
```

❑❷❸❹❺❻❼❽❾❿⓫⓬⓭⓮⓯⓰⓱⓲⓳⓴⓵⓶⓷⓸⓹⓺⓻⓼⓽⓾⓿ MySQL ❶❷❸❹❺❻❼❽❾❿⓫⓬⓭⓮⓯⓰⓱⓲⓳⓴⓵⓶⓷⓸⓹⓺⓻⓼⓽⓾⓿ MySQL ❶❷❸❹❺❻❼❽❾❿⓫⓬⓭⓮⓯⓰⓱⓲⓳⓴⓵⓶⓷⓸⓹⓺⓻⓼⓽⓾⓿ [mysqld]❶❷❸❹❺❻❼❽❾❿⓫⓬⓭⓮⓯⓰⓱⓲⓳⓴⓵⓶⓷⓸⓹⓺⓻⓼⓽⓾⓿

```
[mysqld]  
  
transaction-isolation=READ-COMMITTED
```

❶❷❸❹❺❻❼❽❾❿⓫⓬⓭⓮⓯⓰⓱⓲⓳⓴⓵⓶⓷⓸⓹⓺⓻⓼⓽⓾⓿

```
mysql>SELECT@@tx_isolation\G;  
  
*****1.row*****  
  
@@tx_isolation:REPEATABLE-READ  
  
1 row in set(0.01 sec)
```

❶❷❸❹❺❻❼❽❾❿⓫⓬⓭⓮⓯⓰⓱⓲⓳⓴⓵⓶⓷⓸⓹⓺⓻⓼⓽⓾⓿

```
mysql>SELECT@@global.tx_isolation\G;  
  
*****1.row*****  
  
@@global.tx_isolation:REPEATABLE-READ  
  
1 row in set(0.00 sec)
```

MySQL SERIALIZABLE isolation InnoDB SELECT LOCK IN SHARE MODE SERIALIZABLE well-formed two-phrased

InnoDB REPEATABLE READ 3° SERIALIZABLE SERIALIZABLE InnoDB

READ COMMITTED gap lock InnoDB gap lock MySQL 5.1 READ COMMITTED replication ROW STATEMENT

```
mysql>CREATE TABLE a(
-b INT,PRIMARY KEY(b)
-)ENGINE=INNODB;

Query OK,0 rows affected(0.01 sec)

mysql>SET@@tx_isolation='READ-COMMITTED';

Query OK,0 rows affected(0.00 sec)

mysql>SELECT@@tx_isolation\G;

*****1.row*****

@@tx_isolation:REPEATABLE-READ

1 row in set(0.00 sec)

mysql>BEGIN;

Query OK,0 rows affected(0.00 sec)

mysql>INSERT INTO a SELECT 1;

ERROR 1598(HY000):Binary logging not possible.Message:Transaction level'Read-Committed'in InnoDB is not safe for binlog
mode'Statement'
```

MySQL 5.0innodb_locks_unsafe_for_binlog1STATEMENT READ COMMITTED

```
mysql>SELECT @@version\G
```

```
*****1.row*****
```

```
@@version:5.0.77-log
```

```
1 row in set(0.00 sec)
```

```
mysql>SHOW VARIABLES LIKE 'innodb_locks_unsafe_for_binlog'\G;
```

```
*****1.row*****
```

```
Variable_name:innodb_locks_unsafe_for_binlog
```

```
Value:ON
```

```
1 row in set(0.00 sec)
```

```
mysql>SET @@tx_isolation='READ-COMMITTED';
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>BEGIN;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>INSERT INTO a SELECT 1;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>COMMIT;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>SELECT*FROM a\G;
```

```
*****1.row*****
```

```
b:1
```

```
*****2.row*****
```

```
b:2
```

```
*****3.row*****
```

```
b:4
```

```
*****4.row*****
```

```
b:5
```

4 rows in set(0.00 sec)

mysql> master> A

#Session A on master

mysql> BEGIN;

Query OK, 0 rows affected(0.00 sec)

mysql> DELETE FROM a WHERE b=5;

Query OK, 4 rows affected(0.01 sec)

mysql> master> B

#Session B on master

mysql> BEGIN;

Query OK, 0 rows affected(0.00 sec)

mysql> INSERT INTO a SELECT 3;

Query OK, 0 rows affected(0.01 sec)

mysql> COMMIT;

Query OK, 0 rows affected(0.00 sec)

mysql> A a

#Session A on master

mysql> COMMIT;

Query OK, 0 rows affected(0.00 sec)

mysql> SELECT * FROM a\G;

*****1. row*****

b:3

slave

#Slave

mysql[SELECT*FROM a;

Empty set(0.00 sec)

READ COMMITTED gap lock
B5

STATEMENT master SQL master
STATEMENT

READ
REPEATABLE master slave

MySQL 5.1 ROW
READ COMMITTED READ
COMMITTED ROW
SQL
InnoDB HeikkiTuuri
<http://bugs.mysql.com/bug.php?id=33210> ROW

7.7 数据库

7.7.1 MySQL数据库

InnoDB数据库支持XA事务，XA事务是分布式事务，它支持跨数据库的事务。InnoDB数据库支持transactional resources，它支持ACID事务。InnoDB数据库支持SERIALIZABLE。

XA事务是分布式事务，它支持跨数据库的事务。MySQL数据库支持Oracle数据库，SQL Server数据库支持XA事务。David数据库支持10 000事务，Mariah数据库支持。

#Bank@Shanghai

UPDATE account SET money=money-10000 WHERE user='David';

#Bank@Beijing

UPDATE account SET money=money+10000 WHERE user='Mariah';

David数据库支持Mariah数据库，David数据库支持Mariah数据库。

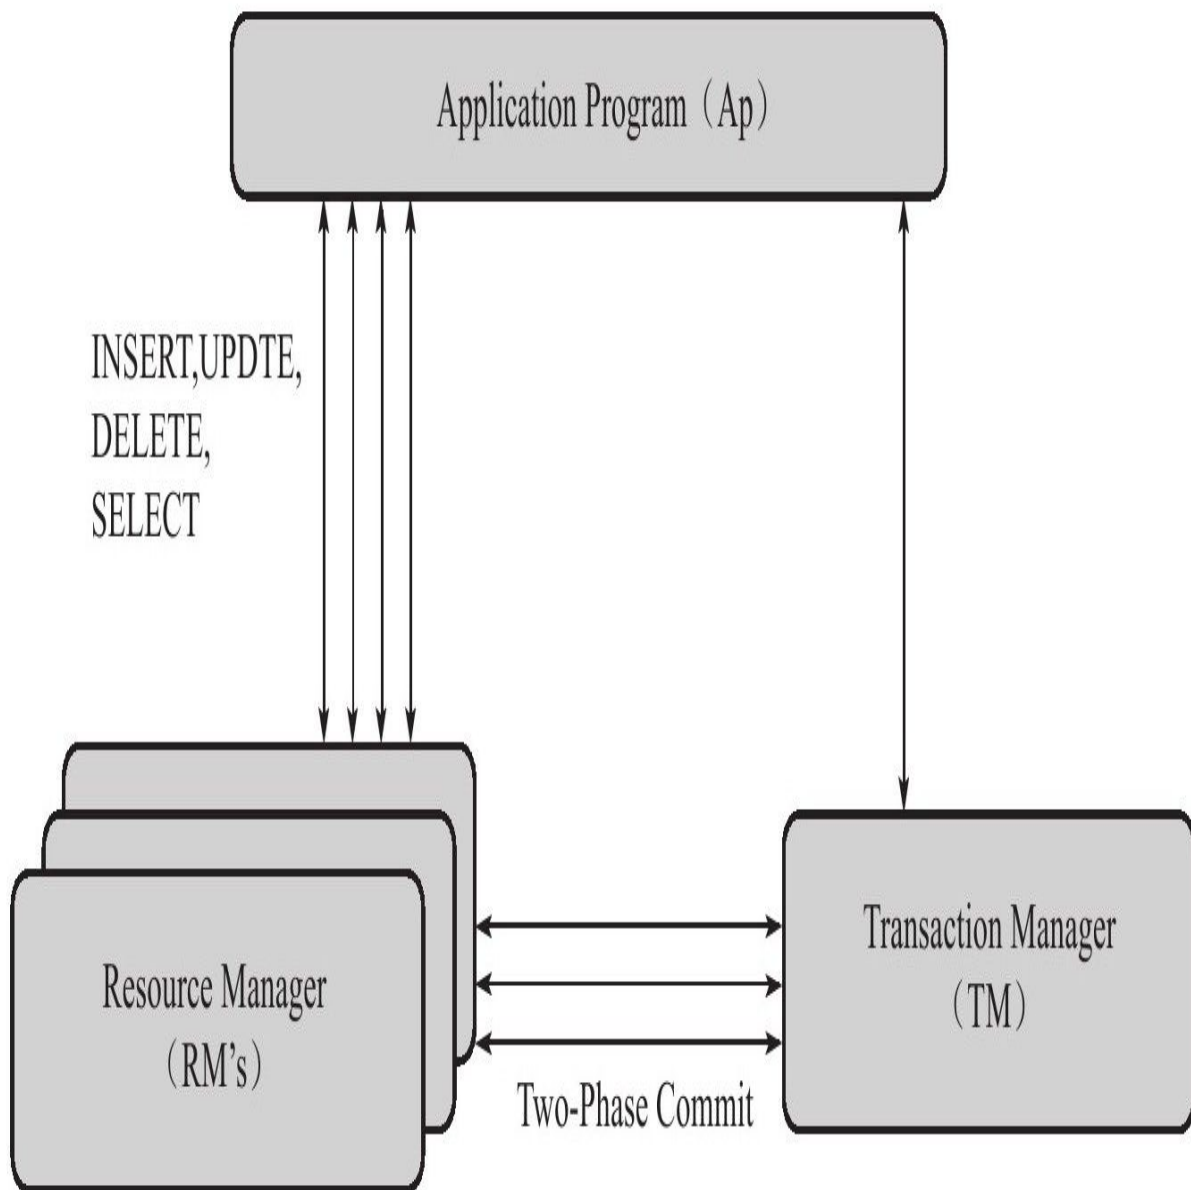
XA事务是Resource Managers，Transaction Manager，Application Program。

数据库支持事务。

数据库支持事务。

数据库支持事务。

MySQLMySQLMySQL
7-22



7-22

two-phase commit
PREPARE
ROLLBACK COMMIT

PREPARE COMMIT ROLLBACK

MySQLXA SQL

XA{START|BEGIN}xid[JOIN|RESUME]

XA END xid[SUSPEND[FOR MIGRATE]]

XA PREPARE xid

XA COMMIT xid[ONE PHASE]

XA ROLLBACK xid

XA RECOVER

XA

mysqlXA START'a';

Query OK,0 rows affected(0.00 sec)

mysqlINSERT INTO z SELECT 11;

Query OK,1 row affected(0.00 sec)

Records:1 Duplicates:0 Warnings:0

mysqlXA END'a';

Query OK,0 rows affected(0.00 sec)

mysqlXA PREPARE'a';

Query OK,0 rows affected(0.05 sec)

mysqlXA RECOVER\G;

*****1.row*****

formatID:1

gtrid_length:1

bqual_length:0

data:a

1 row in set(0.00 sec)

```
mysql>XA COMMIT'a';
```

Query OK, 0 rows affected (0.05 sec)

MySQL
Java JTA Java
Transaction API
MySQL
API JTA MySQL

```
import java.sql.Connection;

import javax.sql.XAConnection;

import javax.transaction.xa.*;

import com.mysql.jdbc.jdbc2.optional.MysqlXADataSource;

import java.sql.*;

class MyXid implements Xid

{

    public int formatId;

    public byte gtrid[];

    public byte bqual[];

    public MyXid(){

    }

    public MyXid(int formatId,byte gtrid[],byte bqual[])

    {

        this.formatId=formatId;

        this.gtrid=gtrid;

        this.bqual=bqual;

    }

    public int getFormatId()

    {

        return formatId;

    }

}
```

```

public byte[]getBranchQualifier()

{

return bqual;

}

public byte[]getGlobalTransactionId()

{

return gtrid;

}

}

public class xa_demo{

public static MysqlXADataSource GetDataSource(

String connString,

String user,

String passwd){

try{

MysqlXADataSource ds=new MysqlXADataSource();

ds.setUrl(connString);

ds.setUser(user);

ds.setPassword(passwd);

return ds;

}

catch(Exception e){

System.out.println(e.toString());

return null;

}

}

public static void main(String[]args){

String connString1="jdbc:mysql://192.168.24.43:3306/bank_shanghai";

String connString2="jdbc:mysql://192.168.24.166:3306/bank_

beijing";

```

```

try{

MysqlXADatasource ds1=

GetDataSource(connString1,"peter","12345");

MysqlXADatasource ds2=

GetDataSource(connString2,"david","12345");

XAConnection xaConn1=ds1.getXAConnection();

XAResource xaRes1=xaConn1.getXAResource();

Connection conn1=xaConn1.getConnection();

Statement stmt1=conn1.createStatement();

XAConnection xaConn2=ds2.getXAConnection();

XAResource xaRes2=xaConn2.getXAResource();

Connection conn2=xaConn2.getConnection();

Statement stmt2=conn2.createStatement();

Xid xid1=new MyXid(

100,

new byte[]{0x01},

new byte[]{0x02});

Xid xid2=new MyXid(

100,

new byte[]{0x11},

new byte[]{0x12});

try{

xaRes1.start(xid1,XAResource.TMNOFLAGS);

stmt1.execute("

UPDATE account SET money=money-10000

WHERE user='david' "

);

xaRes1.end(xid1,XAResource.TMSUCCESS);

xaRes2.start(xid2,XAResource.TMNOFLAGS);

stmt2.execute("

```

```
UPDATE account SET money=money+10000
```

```
WHERE user='mariah'"
```

```
);
```

```
xaRes2.end(xid2,XAResource.TMSUCCESS);
```

```
int ret2=xaRes2.prepare(xid2);
```

```
int ret1=xaRes1.prepare(xid1);
```

```
if(ret1==XAResource.XA_OK
```

```
    &&ret2==XAResource.XA_OK){
```

```
    xaRes1.commit(xid1,false);
```

```
    xaRes2.commit(xid2,false);
```

```
}
```

```
}catch(Exception e){
```

```
    e.printStackTrace();
```

```
}
```

```
}catch(Exception e){
```

```
    System.out.println(e.toString());
```

```
}
```

```
}
```

```
}
```

```
innodb_support_xaXAON
```

```
mysql>SHOW VARIABLES LIKE'innodb_support_xa'\G;
```

```
*****1. row*****
```

```
Variable_name:innodb_support_xa
```

```
Value:ON
```

```
1 row in set(0.01 sec)
```

7.7.2 两阶段XA

两阶段XA是MySQL中实现分布式事务的一种方式。MySQL通过两阶段XA协议，可以在多个MySQL实例之间实现事务的复制和一致性。

两阶段XA协议分为两个阶段：第一阶段是准备阶段，第二阶段是提交阶段。在准备阶段，MySQL会记录binlog和InnoDB redo log。在提交阶段，MySQL会将binlog复制到slave的relay log，并写入InnoDB redo log。最后，MySQL会将slave的redo log复制到master的redo log，完成事务的复制。

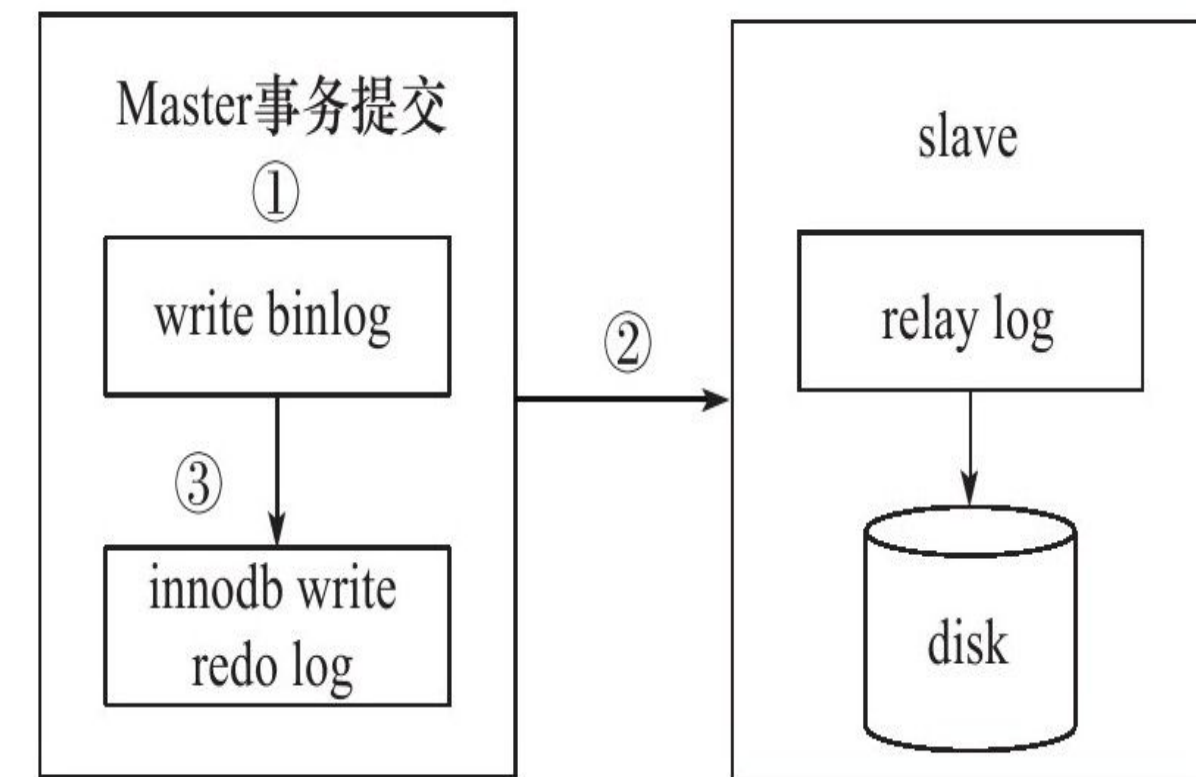


图 7-23 两阶段XA replication流程图

图7-23展示了两阶段XA复制的流程图。① MySQL主节点记录binlog。② MySQL主节点将binlog复制到slave的relay log。③ MySQL主节点将binlog复制到InnoDB redo log。最后，MySQL主节点将slave的redo log复制到master的redo log，完成事务的复制。

图7-24 MySQL InnoDB 引擎的 XA 事务提交过程

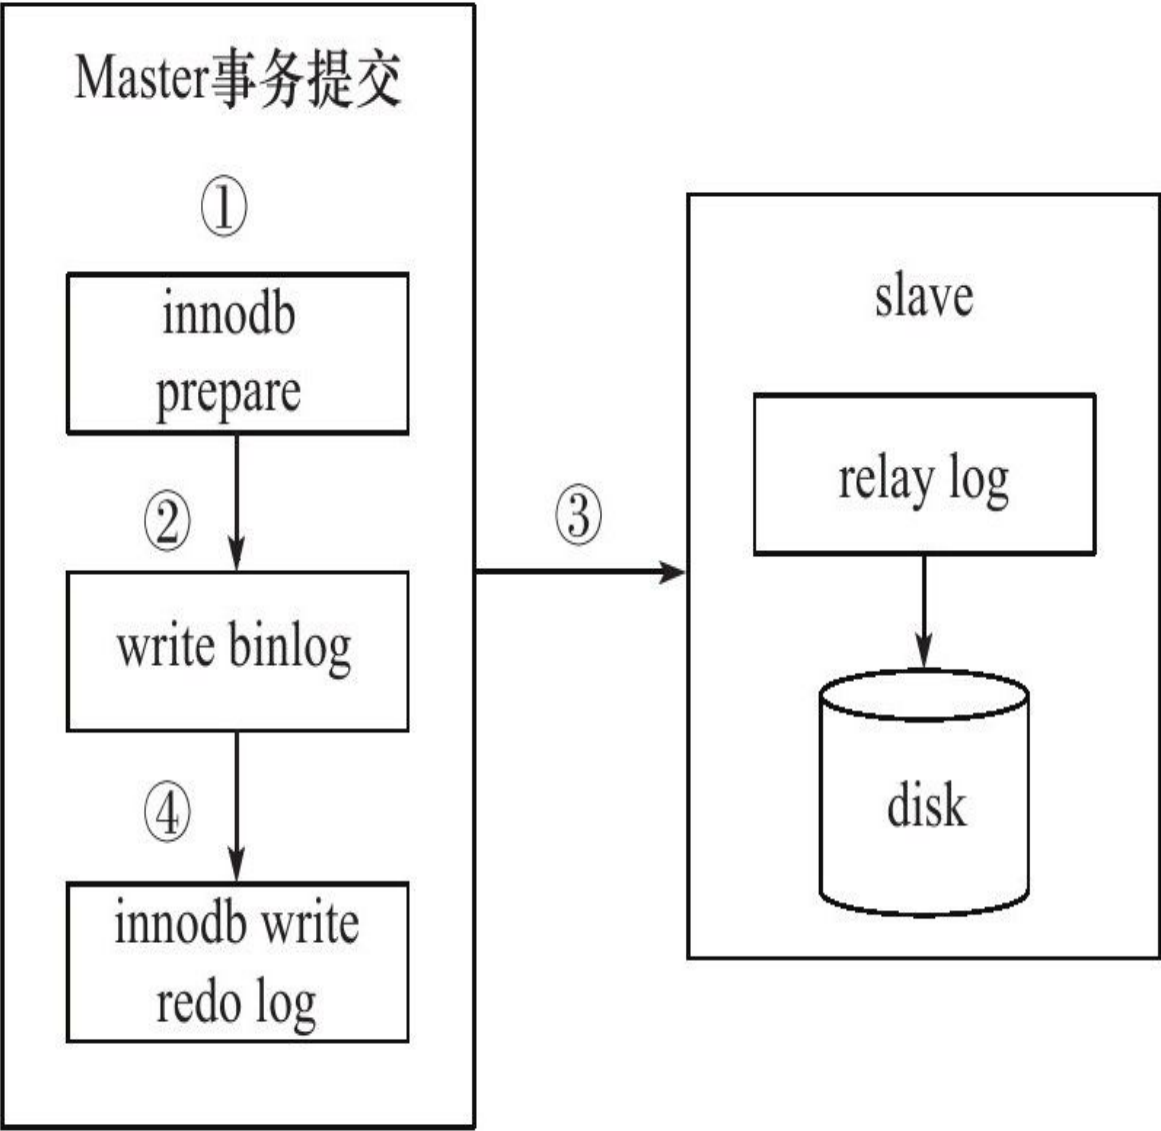


图 7-24 MySQL 引擎的 XA 事务提交过程

7.8 数据库事务

7.8.1 数据库事务

数据库事务是指一个完整的工作单元，事务中的操作要么全部成功，要么全部失败。

```
CREATE PROCEDURE load1(count INT UNSIGNED)

BEGIN

DECLARE s INT UNSIGNED DEFAULT 1;

DECLARE c CHAR(80)DEFAULT REPEAT('a',80);

WHILE s<=count DO

INSERT INTO t1 SELECT NULL,c;

COMMIT;

SET s=s+1;

END WHILE;

END;
```

数据库事务是指一个完整的工作单元，事务中的操作要么全部成功，要么全部失败。

数据库事务是指一个完整的工作单元，事务中的操作要么全部成功，要么全部失败。

```
CREATE PROCEDURE load2(count INT UNSIGNED)

BEGIN

DECLARE s INT UNSIGNED DEFAULT 1;

DECLARE c CHAR(80)DEFAULT REPEAT('a',80);

WHILE s<=count DO

INSERT INTO t1 SELECT NULL,c;

SET s=s+1;

END WHILE;
```



```
END;
```

```
10 00050005000load3
INSERT
```

```
CREATE PROCEDURE load3(count INT UNSIGNED)
```

```
BEGIN
```

```
DECLARE s INT UNSIGNED DEFAULT 1;
```

```
DECLARE c CHAR(80)DEFAULT REPEAT('a',80);
```

```
START TRANSACTION;
```

```
WHILE s<=count DO
```

```
INSERT INTO t1 SELECT NULL,c;
```

```
SET s=s+1;
```

```
END WHILE;
```

```
COMMIT;
```

```
END;
```

```
3
```

```
mysql>CALL load1(10000);
```

```
Query OK,0 rows affected(1 min 3.15 sec)
```

```
mysql>TRUNCATE TABLE t1;
```

```
Query OK,0 rows affected(0.05 sec)
```

```
mysql>CALL load2(10000);
```

```
Query OK,1 row affected(1 min 1.69 sec)
```

```
mysql>TRUNCATE TABLE t1;
```

```
Query OK,0 rows affected(0.05 sec)
```

```
mysql>CALL load3(10000);
```

Query OK, 0 rows affected (0.63 sec)

load1load2
10 000load31
load2load3

mysqlBEGIN;

Query OK, 0 rows affected (0.00 sec)

mysqlCALL load2(10000);

Query OK, 1 row affected (0.56 sec)

mysqlCOMMIT;

Query OK, 0 rows affected (0.03 sec)

InnoDB
OracleundoSnapshot Too
OldMySQLInnoDB

7.8.2 事务控制

数据库事务控制是数据库DBA非常重要的工作。MySQL默认autocommit=1，即自动提交。

```
mysql> SET autocommit=0;
```

```
Query OK, 0 rows affected (0.00 sec)
```

使用START TRANSACTION或BEGIN开始事务，使用COMMIT或ROLLBACK提交事务，使用SET AUTOCOMMIT=1恢复自动提交。

MySQL C API和MySQL Python API都支持SET AUTOCOMMIT=0。MySQL API中，mysql_commit()和mysql_rollback()函数用于提交和回滚事务。

在MySQL中，事务控制是保证数据一致性的关键。通过事务，可以将多个操作作为一个整体来执行，要么全部成功，要么全部失败。这有助于防止数据不一致的情况发生。

7.8.3 事务回滚

InnoDB 数据库引擎支持 HANDLER 函数，用于处理 SQL 异常。在事务中，如果发生异常，可以使用 HANDLER 函数来指定回滚操作。

```
CREATE PROCEDURE sp_auto_rollback_demo()
BEGIN
DECLARE EXIT HANDLER FOR SQLEXCEPTION ROLLBACK;

START TRANSACTION;

INSERT INTO b SELECT 1;

INSERT INTO b SELECT 2;

INSERT INTO b SELECT 1;

INSERT INTO b SELECT 3;

COMMIT;

END;
```

调用 sp_auto_rollback_demo 存储过程，使用 exit 函数来指定 HANDLER 函数。

```
mysql> SHOW CREATE TABLE b\G;

*****1. row*****

Table: b
Create Table: CREATE TABLE 'b' (
  'a' int(11) NOT NULL DEFAULT '0',
  PRIMARY KEY ('a')
) ENGINE=InnoDB DEFAULT CHARSET=latin1

1 row in set (0.00 sec)
```

mysql>CALL sp_auto_rollback_demo;

Query OK,0 rows affected(0.06 sec)

mysql>SELECT*FROM b;

Empty set(0.00 sec)

```
CREATE PROCEDURE sp_auto_rollback_demo()
BEGIN
  DECLARE EXIT HANDLER FOR SQLEXCEPTION BEGIN ROLLBACK;SELECT-1;END;

  START TRANSACTION;

  INSERT INTO b SELECT 1;

  INSERT INTO b SELECT 2;

  INSERT INTO b SELECT 1;

  INSERT INTO b SELECT 3;

  COMMIT;

  SELECT 1;

  END;
```

mysql>CALL sp_auto_rollback_demo()\G;

*****1.row*****

-1:-1

1 row in set(0.04 sec)

mysql>SELECT*FROM b;

[illegible]

```
Microsoft SQL Server
Microsoft SQL Server SET XABORT ON
Microsoft SQL Server
Microsoft SQL Server MySQL
```

```
MySQLdb.connect(host="192.168.8.7",user="root",passwd="xx",db="test")
```

```
cur=conn.cursor()

cur.execute("SET autocommit=0")

cur.execute("CALL sp_rollback_demo")

cur.execute("COMMIT")

except Exception,e:

cur.execute("ROLLBACK")

print e
```

test_demo.py

```
[root@nineyou0-43]#python test_demo.py

starting rollback

(1062,"Duplicate entry'1'for key'PRIMARY'")
```

1062 Duplicate entry'1'for key'PRIMARY'

7.9 长事务

长事务(Long-Lived Transactions)是指那些在数据库系统中运行时间非常长的事务。例如，一个事务可能需要运行1天，甚至更长。在这种情况下，事务的日志记录可能会占用大量的存储空间，并且可能会影响数据库的性能。因此，数据库系统通常会采用一些策略来管理长事务，例如定期提交事务、使用临时表等。

```
UPDATE account
```

```
SET account_total=account_total+(1+interest_rate)
```

在数据库系统中，长事务可能会导致一些问题。例如，长事务可能会占用大量的存储空间，并且可能会影响数据库的性能。因此，数据库系统通常会采用一些策略来管理长事务，例如定期提交事务、使用临时表等。在ACID原则中，长事务可能会导致一致性问题，因此数据库系统通常会采用一些策略来保证长事务的一致性。

在数据库系统中，长事务可能会导致一些问题。例如，长事务可能会占用大量的存储空间，并且可能会影响数据库的性能。因此，数据库系统通常会采用一些策略来管理长事务，例如定期提交事务、使用临时表等。

```
void ComputeInterest(double interest_rate){  
  
    long last_account_done,max_account_no,log_size;  
  
    int batch_size=100000;  
  
    EXEC SQL SELECT COUNT(*) INTO log_size FROM batchcontext;  
  
    if(SQLCODE!=0||log_size==0){  
  
        EXEC SQL DROP TABLE IF EXISTS batchcontext;  
  
        EXEC SQL CREATE TABLE batchcontext(last_account_done BIGINT);  
  
        last_account_done=0;  
  
        INSERT INTO batchcontext SELECT 0;  
  
    }  
  
    else{  
  
        EXEC SQL SELECT last_account_no
```


7.10 总结

本文介绍了InnoDB的ACID特性，包括Atomicity、Consistency、Isolation和Durability。InnoDB通过redo和undo日志实现事务的持久性和一致性。

本文还介绍了InnoDB的REPEATABLE READ隔离级别，以及SQL中如何设置事务隔离级别。InnoDB的REPEATABLE READ隔离级别可以保证在同一个事务中多次读取同一数据的结果是一致的。

本文最后介绍了SQL中如何设置事务隔离级别，以及MySQL中如何设置事务隔离级别。MySQL的默认事务隔离级别是REPEATABLE READ，与InnoDB的默认事务隔离级别一致。

8

```

DBA
RAID
DBA

```

8.1 备份策略

备份策略是指数据库备份的方法和策略。

☐ Hot Backup策略

☐ Cold Backup策略

☐ Warm Backup策略

Hot Backup策略是指数据库在运行状态下进行备份，MySQL Online Backup策略属于Hot Backup策略。Cold Backup策略是指数据库在关闭状态下进行备份，MySQL Offline Backup策略属于Cold Backup策略。Warm Backup策略是指数据库在运行状态下进行备份，MySQL Warm Backup策略属于Warm Backup策略。

备份策略的选择取决于数据库的规模、数据类型、备份频率等因素。

☐ 备份频率

☐ 备份时间

MySQL备份策略是指数据库备份的方法和策略。SQL备份策略是指数据库备份的方法和策略。mysqldump SELECT*INTO OUTFILE策略是指数据库备份的方法和策略。

ibbackup策略是指数据库备份的方法和策略。xtrabackup策略是指数据库备份的方法和策略。

备份策略的选择取决于数据库的规模、数据类型、备份频率等因素。

☐ 备份频率

☐ 备份时间

2008

8.2

```

InnoDBMySQLfrm
*.ibdMySQLmy.cnf

```

[illegible][illegible]

| | | | | | | |
|--|--|--|--|--|--|--|
| | | | | | | |
|--|--|--|--|--|--|--|

[illegible]

MySQL

□ □

SQL

| | | | | | | |
|--|--|--|--|--|--|--|
| | | | | | | |
|--|--|--|--|--|--|--|

```
❑ InnoDB [REDACTED] undo  
[REDACTED]
```

MySQL

8.3 备份

8.3.1 mysqldump

mysqldump 是 Igor Romanenko 开发的 dump 工具，支持 MySQL、MySQL、MySQL、Oracle、Microsoft SQL Server 等数据库。

mysqldump 的基本用法如下：

```
shell$mysqldump[arguments]>file_name
```

备份所有数据库：--all-databases

```
shell$mysqldump--all-databases>dump.sql
```

备份指定数据库：--databases db1 db2 db3

```
shell$mysqldump--databases db1 db2 db3>dump.sql
```

备份指定数据库的单个事务：--single-transaction test

```
[root@xen-server]$mysqldump--single-transaction test>test_backup.sql
```

备份指定数据库的单个事务并导出为 cat 格式：--single-transaction test test_backup.sql cat

```
[root@xen-server]$cat test_backup.sql
```

```
--MySQL dump 10.13 Distrib 5.5.1-m2, for unknown-linux-gnu(x86_64)
```

```
--
```



```
--Host:localhost Database:test

-----

--Server version 5.5.1-m2-log

.....

--

--Table structure for table'a'

--

DROP TABLE IF EXISTS'a';

/*!40101 SET@saved_cs_client=@@character_set_client*/;

/*!40101 SET character_set_client=utf8*/;

CREATE TABLE'a'(

'b'int(11)NOT NULL DEFAULT'0',

PRIMARY KEY('b')

)ENGINE=InnoDB DEFAULT CHARSET=latin1;

/*!40101 SET character_set_client=@saved_cs_client*/;

--

--Dumping data for table'a'

--

LOCK TABLES'a'WRITE;

/*!40000 ALTER TABLE'a'DISABLE KEYS*/;

INSERT INTO'a'VALUES(1),(2),(4),(5);

/*!40000 ALTER TABLE'a'ENABLE KEYS*/;

UNLOCK TABLES;

--

--Table structure for table'z'

--

DROP TABLE IF EXISTS'z';

/*!40101 SET@saved_cs_client=@@character_set_client*/;

/*!40101 SET character_set_client=utf8*/;

CREATE TABLE'z'(
```

```
'a'int(11)DEFAULT NULL

)ENGINE=InnoDB DEFAULT CHARSET=latin1;

/*!40101 SET character_set_client=@saved_cs_client*/;

--

--Dumping data for table'z'

--

LOCK TABLES'z'WRITE;

/*!40000 ALTER TABLE'z'DISABLE KEYS*/;

INSERT INTO'z'VALUES(1),(1);

/*!40000 ALTER TABLE'z'ENABLE KEYS*/;

UNLOCK TABLES;

.....

--Dump completed on 2010-08-03 13:36:17
```

数据库备份的SQL语句
 MySQL数据库的备份
 CREATE TABLE...INSERT...SQL...

mysqldump...mysqldump--help...
 ...--lock-tables...-l...

--single-transaction...START TRANSACTION...
 ...InnoDB...
 ...DDL...DDL...

--lock-tables...-l...MyISAM...
 ...InnoDB...
 ...--single-transaction...--lock-tables...--single-
 transaction...exclusive...MySQL...
 MyISAM...InnoDB...--lock-tables
 ...--lock-tables...
 ...

`--lock-all-tables` -x `--lock-tables`

`--add-drop-database` `CREATE DATABASE` `DROP DATABASE` `--all-databases` `--databases` `CREATE DATABASE`

```
[root@xen-server]#mysqldump--single-transaction--add-drop-database--databases test test_backup.sql
```

```
[root@xen-server]#cat test_backup.sql
```

```
--MySQL dump 10.13 Distrib 5.5.1-m2, for unknown-linux-gnu(x86_64)
```

```
.....
```

```
--
```

```
--Current Database:'test'
```

```
--
```

```
/*!40000 DROP DATABASE IF EXISTS'test'*/;
```

```
CREATE DATABASE/*!32312 IF NOT EXISTS*/'test'/*!40100 DEFAULT CHARACTER SET latin1*/;
```

```
USE'test';
```

```
.....
```

`--master-data[=value]` `replication` `value` `1` `CHANGE MASTER` `value` `2` `CHANGE MASTER` `SQL` `value` `1`

```
[root@xen-server]#mysqldump--single-transaction--add-drop-database--master-data=1--databases test test_backup.sql
```

```
[root@xen-server]#cat test_backup.sql
```

```
--MySQL dump 10.13 Distrib 5.5.1-m2, for unknown-linux-gnu(x86_64)
```

```
--
```

```
--Host:localhost Database:test
```

```
-----
```

```
--Server version 5.5.1-m2-log

.....

--

--Position to start replication or point-in-time recovery from

--

CHANGE MASTER TO MASTER_LOG_FILE='xen-server-bin.000006',MASTER_LOG_POS=8095;
```

`value 2` `CHANGE MASTER`

```
[root@xen-server]#mysqldump--single-transaction--add-drop-database--master-data=2--databases test test_backup.sql
```

```
[root@xen-server]#cat test_backup.sql
```

```
--MySQL dump 10.13 Distrib 5.5.1-m2, for unknown-linux-gnu(x86_64)
```

```
--

--Host:localhost Database:test
```

```
-----

--Server version 5.5.1-m2-log
```

```
.....

--

--Position to start replication or point-in-time recovery from

--

--
```

`--master-data` `--lock-tables` `--single-transaction` `--lock-all-tables`

`--events` `-E`

`--routines` `-R`

□--triggers□□□□□□□□

□--hex-blob□□BINARY□VARBINARY□BLOB□BIT□□□□□□□□□□
□□mysqldump□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□□
□□□□□□□□□□□□--hex-blob□□□□□□□□□□□□□□□□□□□□

```
[root@xen-server]#mysqldump--single-transaction--add-drop-database--master-data=2--no-autocommit--databases test3  
test3_backup.sql
```

```
[root@xen-server]#cat test3_backup.sql
```

```
--MySQL dump 10.13 Distrib 5.5.1-m2, for unknown-linux-gnu(x86_64)
```

```
--
```

```
--Host:localhost Database:test3
```

```
-----
```

```
--Server version 5.5.1-m2-log
```

```
.....
```

```
LOCK TABLES'a'WRITE;
```

```
/*!40000 ALTER TABLE'a'DISABLE KEYS*/;
```

```
setautocommit=0;
```

```
INSERT INTO'a'VALUES(0x610000000000000000000000);
```

```
/*!40000 ALTER TABLE'a'ENABLE KEYS*/;
```

```
UNLOCK TABLES;
```

□□□□□□□□0x61000000000000000000000000□□□□□□□□□□□□□□□□

□--tab=path□-T path□□□□TAB□□□□□□□□□□□□mysqldump□□
□□□□CREATE TABLE□□□table_name.sql□□□□□□□□□□
tbl_name.txt□□□□□□□--fields-terminated-by=...□--fields-
enclosed-by=...□--fields-optionally-enclosed-by=...□--fields-
escaped-by=...□--lines-terminated-by=...□□□□□□□□□□□□□□□□
□□

```
[root@xen-server test]#mysqldump--single-transaction--add-drop-database--tab="/usr/local/mysql/data/test"test
```

```
[root@xen-server test]#ls-lh
```

```
total 244K
```

```
-rw-rw----1 mysql mysql 8.4K Jul 21 16:02 a.frm
```

```
-rw-rw----1 mysql mysql 96K Jul 22 17:18 a.ibd
```

```
-rw-r--r--1 root root 1.3K Aug 3 15:36 a.sql
```

```
-rw-rw-rw-1 mysql mysql 8 Aug 3 15:36 a.txt
```

```
-rw-rw----1 mysql mysql 65 Jul 17 15:54 db.opt
```

```
-rw-rw----1 mysql mysql 8.4K Aug 2 17:22 z.frm
```

```
-rw-rw----1 mysql mysql 96K Aug 2 17:22 z.ibd
```

```
-rw-r--r--1 root root 1.3K Aug 3 15:36 z.sql
```

```
-rw-rw-rw-1 mysql mysql 4 Aug 3 15:36 z.txt
```

```
-----
```

```
--Server version 5.5.1-m2-log
```

```
/*!40101 SET@OLD_CHARACTER_SET_CLIENT=@@CHARACTER_SET_CLIENT*/;
```

```
/*!40101 SET@OLD_CHARACTER_SET_RESULTS=@@CHARACTER_SET_RESULTS*/;
```

```
/*!40101 SET@OLD_COLLATION_CONNECTION=@@COLLATION_CONNECTION*/;
```

```
/*!40101 SET NAMES utf8*/;
```

```
/*!40103 SET@OLD_TIME_ZONE=@@TIME_ZONE*/;
```

```
/*!40103 SET TIME_ZONE='+00:00'*/;
```

```
/*!40101 SET@OLD_SQL_MODE=@@SQL_MODE,SQL_MODE='*/;
```

```
/*!40111 SET@OLD_SQL_NOTES=@@SQL_NOTES,SQL_NOTES=0*/;
```

```
--
```

```
--Table structure for table'a'
```

```
--
```

```
DROP TABLE IF EXISTS'a';
```

```
/*!40101 SET@saved_cs_client=@@character_set_client*/;
```

```
/*!40101 SET character_set_client=utf8*/;
```

```
CREATE TABLE'a'(  
  'b'int(11)NOT NULL DEFAULT'0',  
  PRIMARY KEY('b')
```

```

)ENGINE=InnoDB DEFAULT CHARSET=latin1;

/*!40101 SET character_set_client=@saved_cs_client*/;

/*!40103 SET TIME_ZONE=@OLD_TIME_ZONE*/;

/*!40101 SET SQL_MODE=@OLD_SQL_MODE*/;

/*!40101 SET CHARACTER_SET_CLIENT=@OLD_CHARACTER_SET_CLIENT*/;

/*!40101 SET CHARACTER_SET_RESULTS=@OLD_CHARACTER_SET_RESULTS*/;

/*!40101 SET COLLATION_CONNECTION=@OLD_COLLATION_CONNECTION*/;

/*!40111 SET SQL_NOTES=@OLD_SQL_NOTES*/;

--Dump completed on 2010-08-03 15:36:56

[root@xen-server test]#cat a.txt

```

```

1
2
4
5

```

数据库DBA
 SELECT...INTO OUTFILE
 mysqldump

--where='where_condition' -w'where_condition'
 b a a 2

```

[root@xen-server bin]#mysqldump--single-transaction--where='b2'test a.a.sql

```

```

[root@xen-server bin]#cat a.sql

```

```

--MySQL dump 10.13 Distrib 5.5.1-m2, for unknown-linux-gnu(x86_64)

```

```

--

```

```

--Host:localhost Database:test

```

```

-----

```

```

--Server version 5.5.1-m2-log

```

```

.....

```

```

--

```

```

--Dumping data for table'a'

```

--

--WHERE:b□2

LOCK TABLES'a'WRITE;

/*!40000 ALTER TABLE'a'DISABLE KEYS*/;

INSERT INTO'a'VALUES(4),(5);

/*!40000 ALTER TABLE'a'ENABLE KEYS*/;

UNLOCK TABLES;

/*!40103 SET TIME_ZONE=@OLD_TIME_ZONE*/;

.....

8.3.2 SELECT...INTO OUTFILE

SELECT...INTO
SELECT...INTO

```
SELECT[column 1],[column 2]...
```

```
INTO
```

```
OUTFILE'file_name'
```

```
[{FIELDS|COLUMNS}
```

```
[TERMINATED BY'string']
```

```
[[OPTIONALLY]ENCLOSED BY'char']
```

```
[ESCAPED BY'char']
```

```
]
```

```
[LINES
```

```
[STARTING BY'string']
```

```
[TERMINATED BY'string']
```

```
]
```

```
FROM TABLE WHERE.....
```

FIELDS[TERMINATED BY'string']
[[OPTIONALLY]ENCLOSED BY'char']
[ESCAPED BY'char']
[STARTING BY'string']
TERMINATED BY'string'
FIELDS
LINES

```
FIELDS TERMINATED BY'\t'ENCLOSED BY''ESCAPED BY'\\'
```

```
LINES TERMINATED BY'\n'STARTING BY''
```

file_name
mysql:mysql
MySQL

Windows“\r\n”LINES
TERMINATED BY

```
[root@xen-servermysql]#mysql test-e"select*into outfile'/home/mysql/a.txt'fields terminated by','lines terminated by'\r\n'from a";

[root@xen-servermysql]#od-c a.txt

0000000 1,a\r\n 2,b\r\n 3,c\r\n

0000017
```

8.3.3 数据库备份

mysqldump 数据库 数据库名 SQL 语句 数据库名 数据库名

```
[root@xen-server ~]#mysql-uroot-p test_backup.sql
```

Enter password:

数据库 数据库名 SQL 语句 数据库名 数据库名

```
mysql>drop database test;
```

```
ERROR 1010(HY000):Error dropping database(can't rmdir'./test',errno:39)
```

数据库 数据库名 SQL 语句 SOURCE 数据库名 数据库名

```
mysql>source/home/mysql/test_backup.sql;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
Query OK,0 rows affected(0.00 sec)
```

.....

```
Query OK,0 rows affected(0.00 sec)
```

```
Query OK,0 rows affected(0.00 sec)
```

mysqldump 数据库 数据库名 mysqldump 数据库名
数据库名 数据库名 数据库名 数据库名 数据库名 数据库名
mysqldump 数据库名 数据库名 frm 数据库名
数据库名 mysqldump 数据库名

8.3.4 LOAD DATA INFILE

[illegible]

```
LOAD DATA INTO TABLE a IGNORE 1 LINES INFILE '/home/mysql/a.txt'
```

[REPLACE | IGNORE]

INTO TABLE tbl_name

[CHARACTER SET charset_name]

[{FIELDS| COLUMNS}

```
[TERMINATED BY'string']
```

[[OPTIONALLY]ENCLOSED BY'char']

[ESCAPED BY 'char']

]

[LINES

```
[STARTING BY'string']
```

```
[TERMINATED BY'string']
```

]

[IGNORE number LINES]

```
[(col_name_or_user_var,...)]
```

[SET col_name=expr,...]

```

LOAD DATA INFILE 'FILE'
SELECT INTO OUTFILE 'IGNORE number LINES'
LOAD DATA INFILE

```

```
mysql\load data infile'/home/mysql/a.txt'into table a;
```

Query OK, 3 rows affected (0.00 sec)

Records:3 Deleted:0 Skipped:0 Warnings:0

MySQL InnoDB Foreign Key Checks

```
mysql>SET@@foreign_key_checks=0;
```

```
Query OK,0 rows affected(0.00 sec)
```

```
mysql>LOAD DATA INFILE '/home/mysql/a.txt' INTO TABLE a;
```

```
Query OK,4 rows affected(0.00 sec)
```

```
Records:4 Deleted:0 Skipped:0 Warnings:0
```

```
mysql>SET@@foreign_key_checks=1;
```

```
Query OK,0 rows affected(0.00 sec)
```

MySQL InnoDB Foreign Key Checks

```
mysql>CREATE TABLE b(
```

```
- a INT,
```

```
- b INT,
```

```
- c INT,
```

```
- PRIMARY KEY(a)
```

```
- )ENGINE=InnoDB;
```

```
Query OK,0 rows affected(0.01 sec)
```

```
mysql>LOAD DATA INFILE '/home/mysql/a.txt'
```

```
- INTO TABLE b FIELDS TERMINATED BY','(a,b)
```

```
- SET c=a+b;
```

```
Query OK,4 rows affected(0.01 sec)
```

```
Records:4 Deleted:0 Skipped:0 Warnings:0
```

```
mysql>SELECT*FROM b;
```

```
+---+-----+-----+
```

```
|a|b|c|
```

```
+---+-----+-----+
```

```
|1|2|3|
```

|2|3|5|

|4|5|9|

|5|6|11|

+---+-----+-----+

4 rows in set(0.00 sec)

8.3.5 mysqlimport

mysqlimport MySQL LOAD DATA INFILE LOAD DATA INFILE

```
shell:mysqlimport[options]db_name textfile1[textfile2...]
```

LOAD DATA INFILE mysqlimport --
user-thread
mysqlimport
mysqlimport 2

```
[root@xen-servermysql]#mysqlimport--use-threads=2 test/home/mysql/t.txt/home/mysql/s.txt
```

```
test.s:Records:5000000 Deleted:0 Skipped:0 Warnings:0
```

```
test.t:Records:5000000 Deleted:0 Skipped:0 Warnings:0
```

MySQL

```
mysql[SHOW FULL PROCESSLIST\G;
```

```
*****1.row*****
```

```
Id:46
```

```
User:rep
```

```
Host:www.dao.com:1028
```

```
db:NULL
```

```
Command:Binlog Dump
```

```
Time:37651
```

```
State:Master has sent all binlog to slave;waiting for binlog to be updated
```

```
Info:NULL
```

```
*****2.row*****
```


Id:77

User:root

Host:localhost

db:test

Command:Query

Time:0

State:NULL

Info:show full processlist

*****3.row*****

Id:83

User:root

Host:localhost

db:test

Command:Query

Time:73

State:NULL

Info:LOAD DATA INFILE'/home/mysql/t.txt'INTO TABLE't'IGNORE 0 LINES

*****4.row*****

Id:84

User:root

Host:localhost

db:test

Command:Query

Time:73

State:NULL

Info:LOAD DATA INFILE'/home/mysql/s.txt'INTO TABLE's'IGNORE 0 LINES

4 rows in set(0.00 sec)

mysqlimportLOAD DATA INFILE

8.4 point-in-time replication

point-in-time replication MySQL replication

```
[mysqld]
log-bin=mysql-bin
```

3.2.4 InnoDB InnoDB

```
[mysqld]
log-bin=mysql-bin
sync_binlog=1
innodb_support_xa=1
```

FLUSH LOGS

mysqlbinlog mysqlbinlog

```
shell[mysqlbinlog[options]log_file...
```

binlog.0000001

```
shell[mysqlbinlog binlog.0000001|mysql-uroot-p test
```

mysqlbinlog 命令的常用选项

```
shell$mysqlbinlog binlog.[0-10]*|mysql-u root-p test
```

mysqlbinlog 命令的常用选项 SOURCE 选项

```
shell$mysqlbinlog binlog.000001/tmp/statements.sql
```

```
shell$mysqlbinlog binlog.000002/tmp/statements.sql
```

```
shell$mysql-u root-p-e"source/tmp/statements.sql"
```

--start-position --stop-position 选项

```
shell$mysqlbinlog--start-position=107856 binlog.000001|mysql-u root-p test
```

--start-datetime --stop-datetime 选项

8.5 簡介

8.5.1 ibbackup

ibbackup 支援 InnoDB 儲存引擎的 MyISAM 儲存引擎 InnoDB 儲存引擎 InnoDB 儲存引擎

1 支援 InnoDB 儲存引擎 LSN

2 支援 InnoDB 儲存引擎

3 支援 InnoDB 儲存引擎 LSN

4 支援 InnoDB 儲存引擎

Microsoft SQL Server Oracle 儲存引擎
儲存引擎 儲存引擎 儲存引擎 儲存引擎 儲存引擎 儲存引擎
ibbackup 支援

SQL 儲存引擎

儲存引擎 儲存引擎 儲存引擎 儲存引擎 儲存引擎 儲存引擎

儲存引擎 儲存引擎 儲存引擎 儲存引擎 儲存引擎 儲存引擎

ibbackup 支援 Linux Windows UNIX 儲存引擎

ibbackup 支援 InnoDB 儲存引擎

儲存引擎

儲存引擎

ibbackup 支援 InnoDB 儲存引擎 儲存引擎 儲存引擎 儲存引擎 儲存引擎 儲存引擎
Percona 儲存引擎 儲存引擎 儲存引擎 儲存引擎 儲存引擎 儲存引擎

XtraBackupibbackup
XtraBackup

8.5.2 XtraBackup

XtraBackup is a Percona MySQL 5.0 backup tool.
XtraBackup is GPL v2.
<https://launchpad.net/percona-xtrabackup>

xtrabackup

```
xtrabackup--backup|--prepare[OPTIONS]
```

xtrabackup

(The defaults options should be given as the first argument)

--print-defaults Prints the program's argument list and exit.

--no-defaults Don't read the default options from any file.

--defaults-file=Read the default options from this file.

--defaults-extra-file=Read this file after the global options files have been read.

--target-dir=The destination directory for backups.

--backup Make a backup of a mysql instance.

--stats Calculate the statistic of the datadir(it is recommended you take mysqld offline).

--prepare Prepare a backup so you can start mysql server with your restore.

--export Create files to import to another database after it has been prepared.

--print-param Print the parameters of mysqld that you will need for a forcycopyback.

--use-memory=This value is used instead of buffer_pool_size.

--suspend-at-end Creates a file called xtrabackup_suspended and waits until the user deletes that file at the end of the backup.

--throttle=(use with--backup)Limits the IO operations(pairs of reads and writes)per second to the values set here.

--log-stream outputs the contents of the xtrabackup_logfile to stdout.

--incremental-lsn=(use with--backup)Copy only.ibd pages newer than the specified LSN high:low.##ATTENTION##:checkpoint lsn*must*be used.Be Careful!

--incremental-basedir=(use with--backup)Copy only.ibd pages newer than the existing backup at the specified directory.

--incremental-dir=(use with--prepare)Apply.delta files and logfiles located in the specified directory.

--tables=name Regular Expression list of table names to be backed up.

--create-ib-logfile(NOT CURRENTLY IMPLEMENTED)will create ib_logfile*after a--prepare.

###If you want to create ib_logfile*only re-execute this

command using the same options.###

--datadir=name Path to the database root.

--tmpdir=name Path for temporary files.Several paths may be specified as a colon(:)separated string.

If you specify multiple paths they are used round-robin.

XXXXXXXXXXXXXXXXXXXXXXXXXXXX

#!/xtrabackup--backup

./xtrabackup Ver alpha-0.2 for 5.0.75 unknown-linux-gnu(x86_64)

log scanned up to(0 1009910580)

Copying./ibdata1

to/home/kinoyasu/xtrabackup_work/mysql-5.0.75/innobase/xtrabackup/tmp2/ibdata1

...done

Copying./tpcc/stock.ibd

to/home/kinoyasu/xtrabackup_work/mysql-5.0.75/innobase/xtrabackup/tmp2/tpcc/stock.ibd

...done

Copying./tpcc/new_orders.ibd

to/home/kinoyasu/xtrabackup_work/mysql-5.0.75/innobase/xtrabackup/tmp2/tpcc/new_orders.ibd

...done

Copying./tpcc/history.ibd

to/home/kinoyasu/xtrabackup_work/mysql-5.0.75/innobase/xtrabackup/tmp2/tpcc/history.ibd

...done

Copying./tpcc/customer.ibd

to/home/kinoyasu/xtrabackup_work/mysql-5.0.75/innobase/xtrabackup/tmp2/tpcc/customer.ibd

log scanned up to(0 1010561109)

...done

```
Copying./tpcc/district.ibd

to/home/kinoyasu/xtrabackup_work/mysql-5.0.75/innobase/xtrabackup/tmp2/tpcc/district.ibd

...done

Copying./tpcc/item.ibd

to/home/kinoyasu/xtrabackup_work/mysql-5.0.75/innobase/xtrabackup/tmp2/tpcc/item.ibd

...done

Copying./tpcc/order_line.ibd

to/home/kinoyasu/xtrabackup_work/mysql-5.0.75/innobase/xtrabackup/tmp2/tpcc/order_line.ibd

log scanned up to(0 1012047066)

...done

Copying./tpcc/orders.ibd

to/home/kinoyasu/xtrabackup_work/mysql-5.0.75/innobase/xtrabackup/tmp2/tpcc/orders.ibd

...done

Copying./tpcc/warehouse.ibd

to/home/kinoyasu/xtrabackup_work/mysql-5.0.75/innobase/xtrabackup/tmp2/tpcc/warehouse.ibd

...done

log scanned up to(0 1014592707)

Stopping log copying thread..

Transaction log of lsn(0 1009910580)to(0 1014592707)was copied.
```

```

xtrabackup
1009910580InnoDB
copyCopying...to...
1014592707
```


8.5.3 XtraBackup

MySQL point-in-time XtraBackup InnoDB

1 LSN

2 LSN LSN LSN

XtraBackup

(full backup)

#./xtrabackup--backup--target-dir=/backup/base

...

(incremental backup)

#./xtrabackup--backup--target-dir=/backup/delta--incremental-basedir=/backup/base

...

(prepare)

#./xtrabackup--prepare--target-dir=/backup/base

...

(apply incremental backup)

#./xtrabackup--prepare--target-dir=/backup/base--incremental-dir=/backup/delta

...

/backup/base

/backup/delta

#./xtrabackup--backup

./xtrabackup Ver beta-0.4 for 5.0.75 unknown-linux-gnu(x86_64)

log scanned up to(0 378161500)

...

The latest check point(for incremental):'0:377883685'====LSN

log scanned up to(0 379294296)

Stopping log copying thread..

Transaction log of lsn(0 377883685)to(0 379294296)was copied.

(must do--prepare before the each incremental backup)

#!/xtrabackup--prepare

...

#!/xtrabackup--backup--incremental=0:377883685

incremental backup from 0:377883685 is enabled.

./xtrabackup Ver beta-0.4 for 5.0.75 unknown-linux-gnu(x86_64)

log scanned up to(0 379708047)

Copying./ibdata1

to/home/kinoyasu/xtrabackup_work/mysql-5.0.75/innobase/xtrabackup/tmp_diff/ibdata1.delta

...done

...

The latest check point(for incremental):'0:379438233'====LSN

log scanned up to(0 380663549)

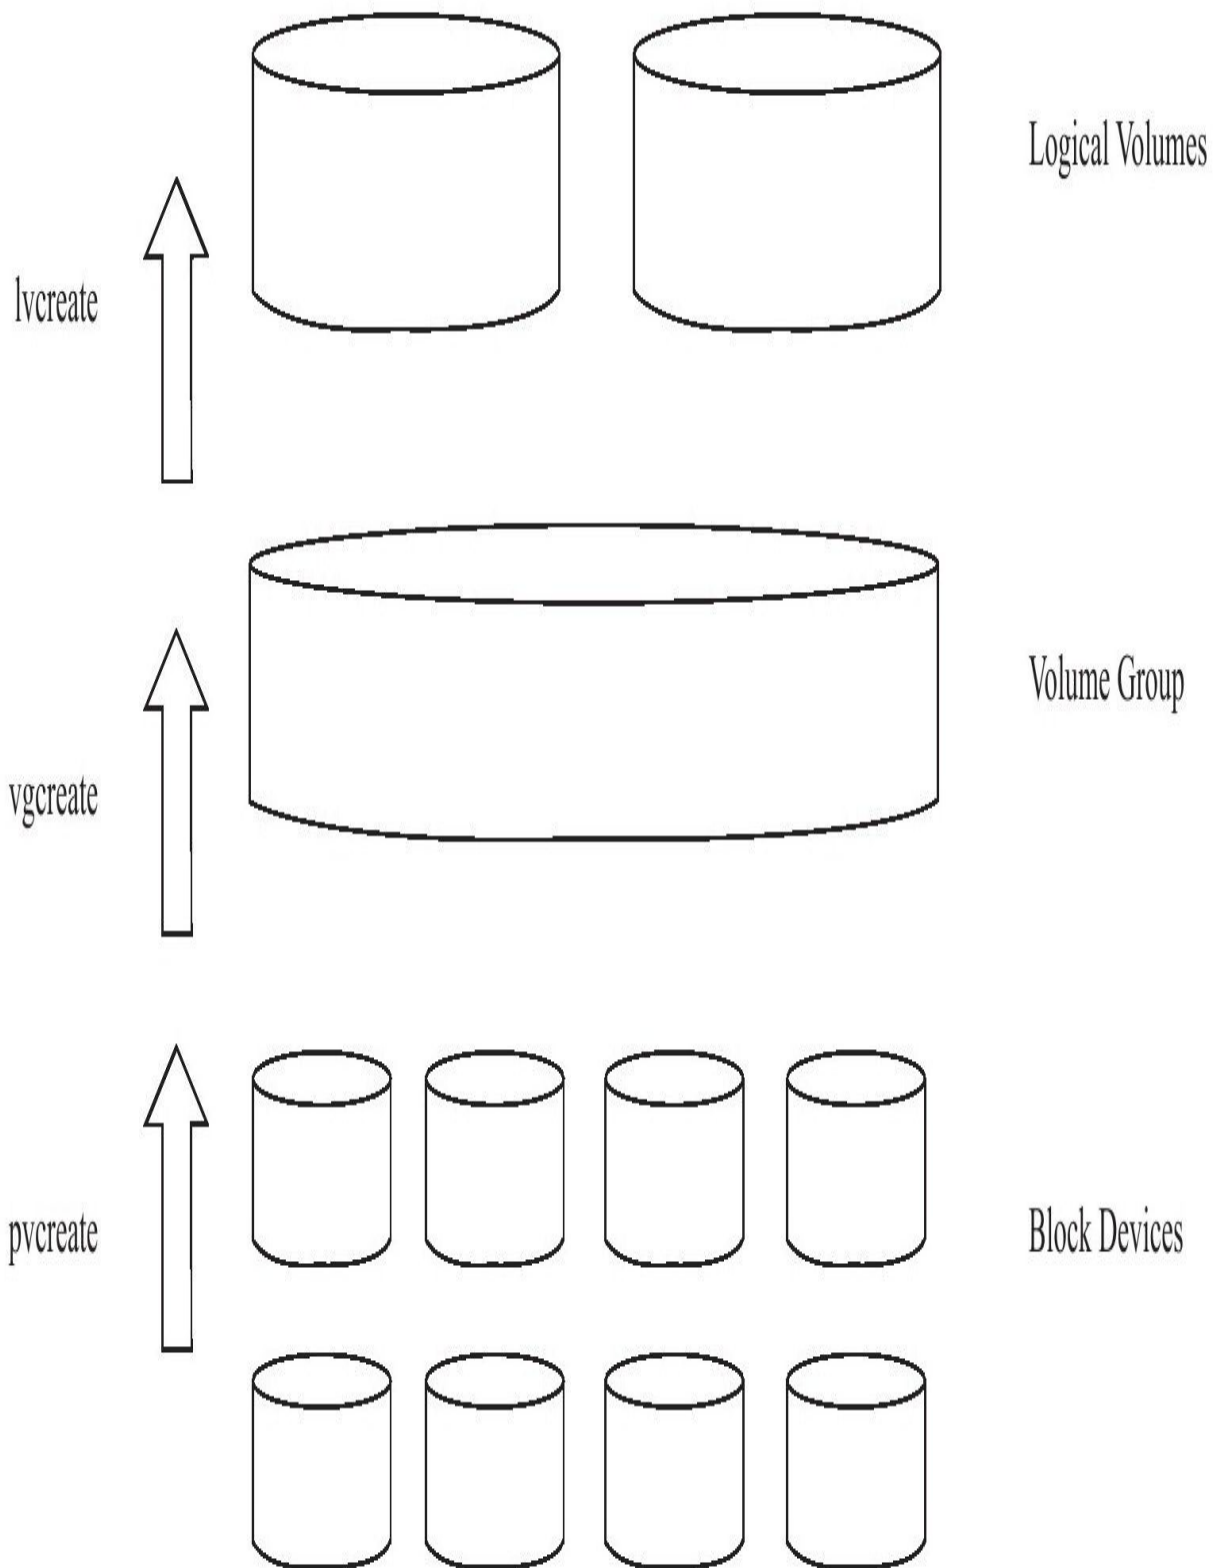
Stopping log copying thread..

Transaction log of lsn(0 379438233)to(0 380663549)was copied.

8.6 □□□□

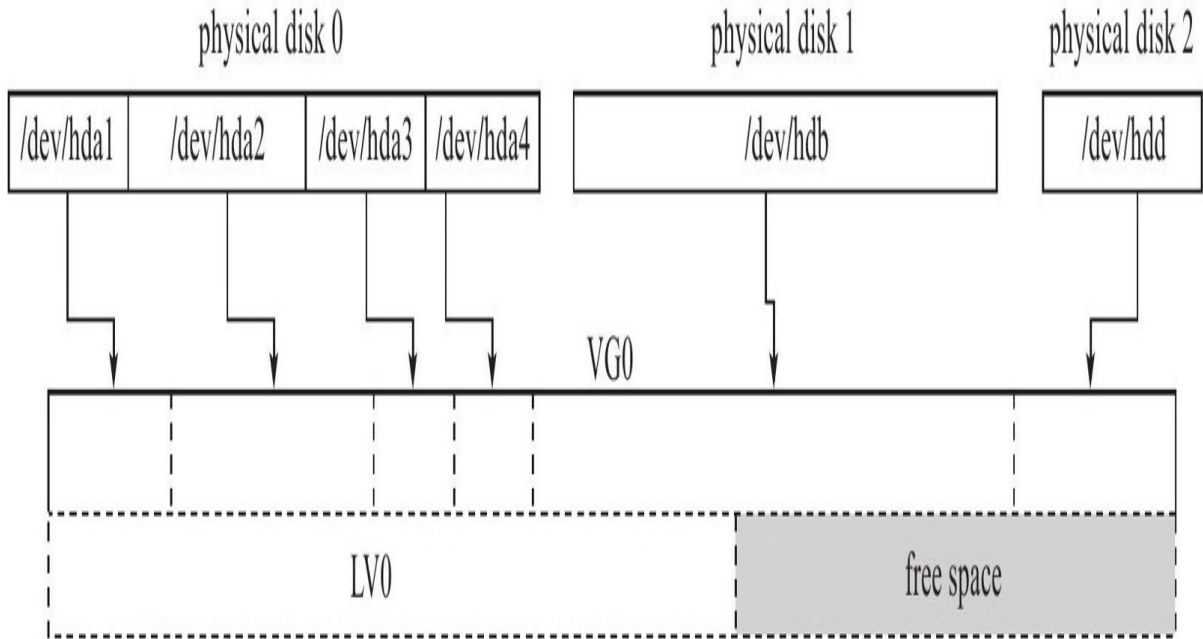
MySQL
FreeBSD UFS
Solaris ZFS
GNU/Linux
Logical Volume Manager LVM
UFS ZFS
LVM

LVM LINUX LVM
LVM
Volume Group Logical
Volumes LVM
LVM
8-1



8-1 LVM

8-2 LVM



8-2

vgdisplay

```
[root@nh124-98]#vgdisplay
```

```
---Volume group---
```

```
VG Name rep
```

```
System ID
```

```
Format lvm2
```

```
Metadata Areas 1
```

```
Metadata Sequence No 1873
```

```
VG Access read/write
```

```
VG Status resizable
```

```
MAX LV 0
```

Cur LV 3
Open LV 1
Max PV 0
Cur PV 1
Act PV 1
VG Size 260.77 GB
PE Size 4.00 MB
Total PE 66758
Alloc PE/Size 66560/260.00 GB
Free PE/Size 198/792.00 MB
VG UUID MQJiye-j4NN-LbZG-F3CQ-UdTU-fo9D-RRfXD5

vgdisplay rep 260.77GB
read/write lvdisk

```
[root@nh124-98]#lvdisplay
---Logical volume---
LV Name/dev/rep/repdata
VG Name rep
LV UUID 7t0lDt-seKZ-ChpY-QMXC-WaFD-zXAl-MRbofK
LV Write Access read/write
LV snapshot status source of
/dev/rep/dho_datasnapshot100805143507[active]
/dev/rep/dho_datasnapshot100805163504[active]
LV Status available
#open 1
LV Size 100.00 GB
Current LE 25600
Segments 1
Allocation inherit
```

Read ahead sectors auto

-currently set to 256

Block device 253:0

---Logical volume---

LV Name/dev/rep/dho_datasnapshot100805143507

VG Name rep

LV UUID fSSXzh-IBnZ-aZIn-eP03-b7pk-CPjN-5xUktE

LV Write Access read only

LV snapshot status active destination for/dev/rep/repdata

LV Status available

#open 0

LV Size 100.00 GB

Current LE 25600

COW-table size 80.00 GB

COW-table LE 20480

Allocated to snapshot 0.13%

Snapshot chunk size 4.00 KB

Segments 1

Allocation inherit

Read ahead sectors auto

-currently set to 256

Block device 253:1

---Logical volume---

LV Name/dev/rep/dho_datasnapshot100805163504

VG Name rep

LV UUID 3B9NP1-qWVG-pfJY-Bdgm-DIdD-dUMu-s2L6qJ

LV Write Access read only

LV snapshot status active destination for/dev/rep/repdata

LV Status available

#open 0

LV Size 100.00 GB

Current LE 25600

COW-table size 80.00 GB

COW-table LE 20480

Allocated to snapshot 0.02%

Snapshot chunk size 4.00 KB

Segments 1

Allocation inherit

Read ahead sectors auto

-currently set to 256

Block device 253:4

3rep
100GB/dev/rep/repdata

LVM Copy-on-write
meta data

8-3LVM B A
C D B

数据来源卷

快照区域

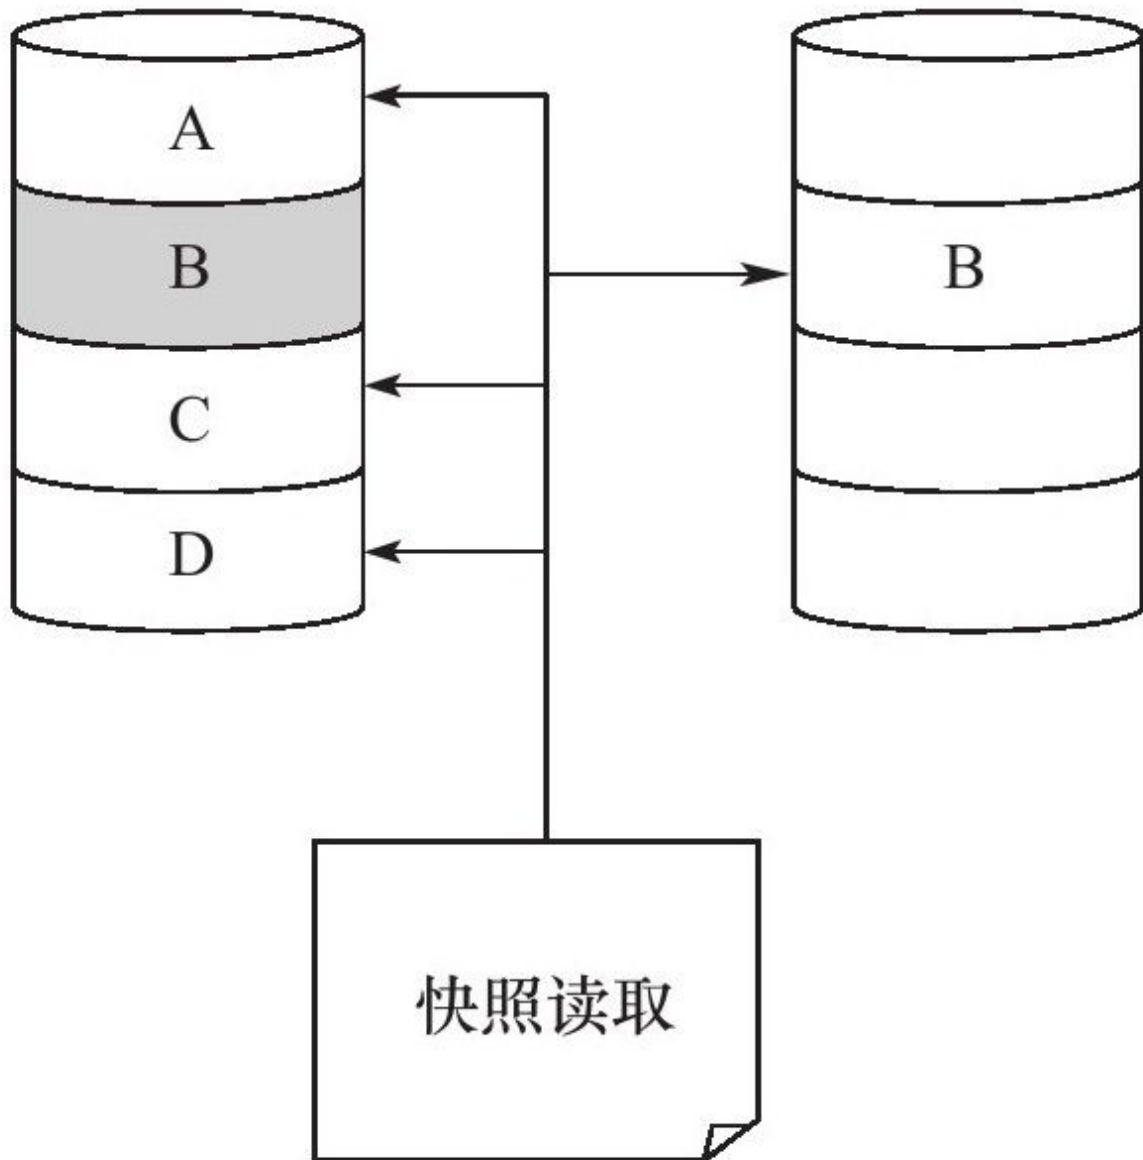


图 8-3 LVM快照

`lvcreate --size 100G --snapshot --permission r-n datasnapshot/dev/rep/repdata`

```
[root@nh119-215 data]#lvcreate--size 100G--snapshot--permission r-n datasnapshot/dev/rep/repdata
```

Logical volume "datasnapshot" created

lvdisplay COW-table size
Allocated to snapshot

[root@nh124-98]#lvdisplay

.....

---Logical volume---

LV Name/dev/rep/dho_datasnapshot100805163504

VG Name rep

LV UUID 3B9NP1-qWVG-pfJY-Bdgm-DIdD-dUMu-s2L6qJ

LV Write Access read only

LV snapshot status active destination for/dev/rep/repdata

LV Status available

#open 0

LV Size 100.00 GB

Current LE 25600

COW-table size 80.00 GB

COW-table LE 20480

Allocated to snapshot 0.04%

Snapshot chunk size 4.00 KB

Segments 1

Allocation inherit

Read ahead sectors auto

-currently set to 256

Block device 253:4

0.04%
Allocated to snapshot

LVM InnoDB InnoDB

```

InnoDB
InnoDB

```

8.7 复制

8.7.1 复制原理

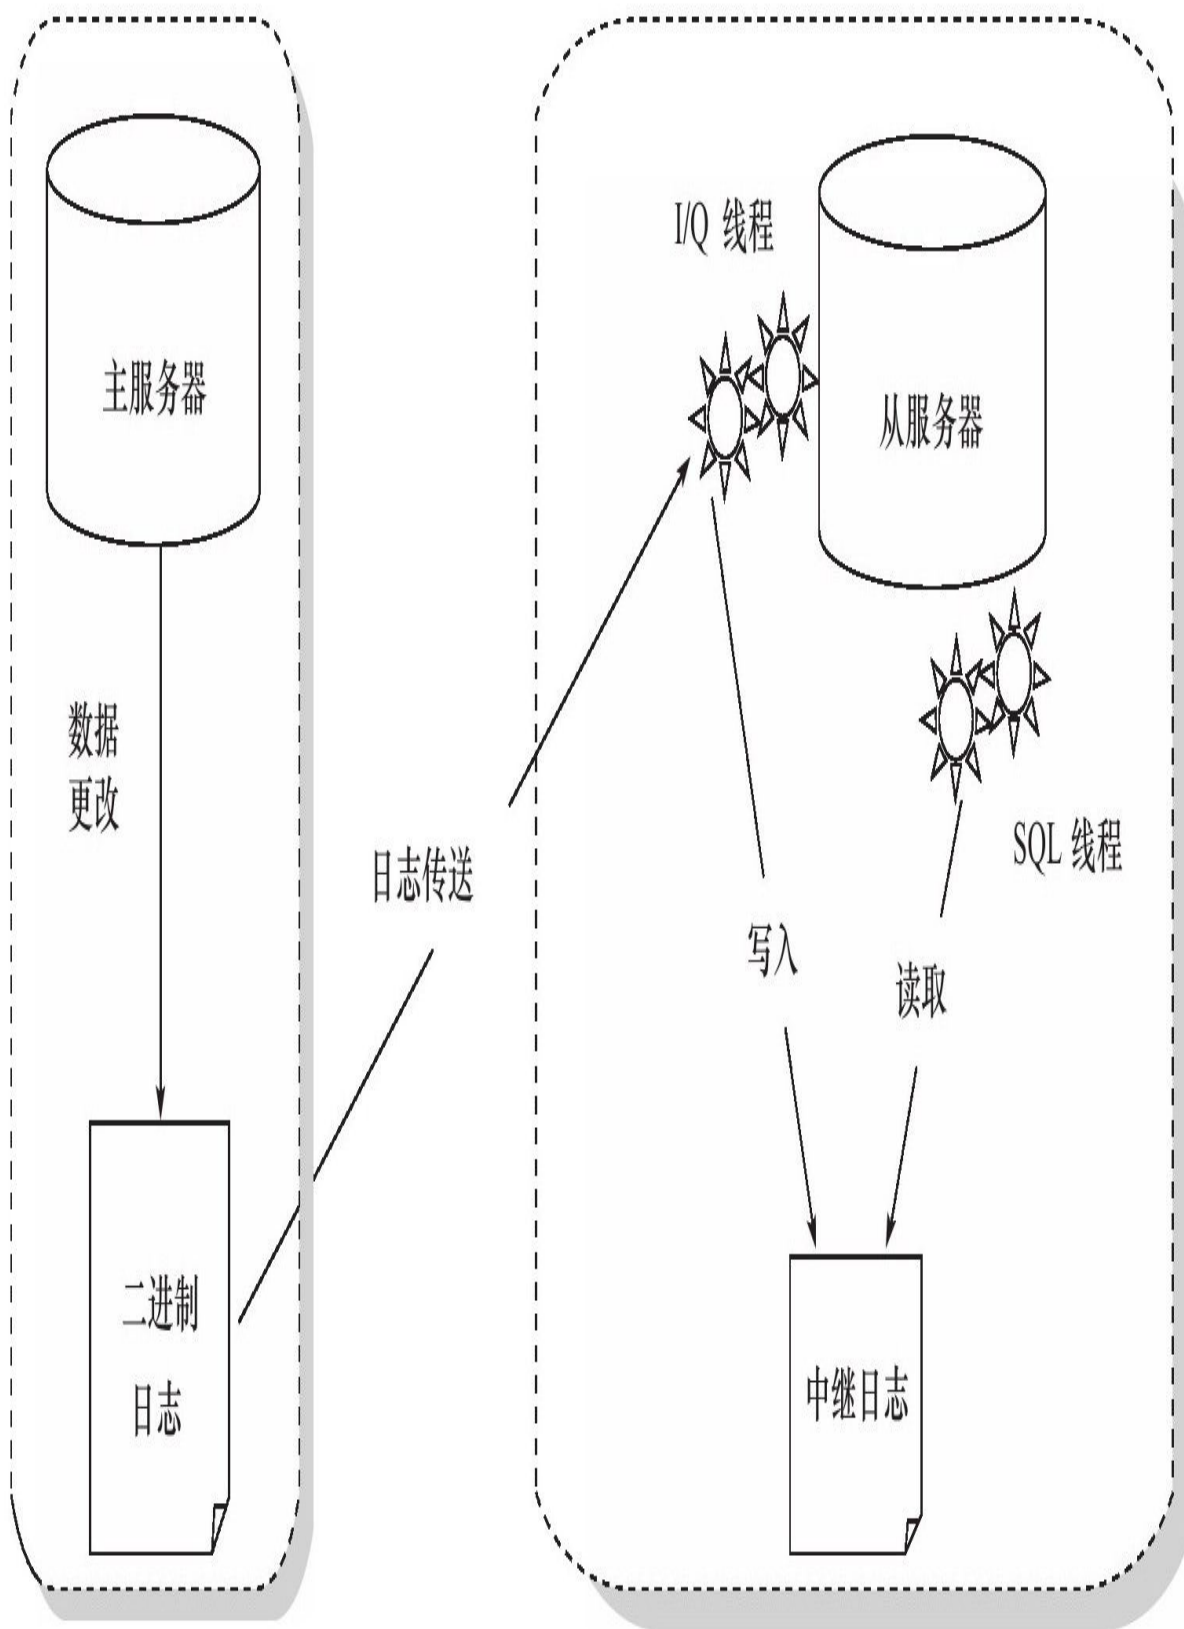
MySQL replication 是 MySQL 数据库的一种功能，它允许将一个数据库的数据复制到另一个数据库。复制原理可以分为 3 个步骤：

1. 主数据库（master）记录所有数据变更到二进制日志（binlog）中。

2. 从数据库（slave）定期从主数据库的 binlog 中读取数据。

3. 从数据库将读取到的数据应用到自己的数据库中。

复制原理的示意图如下所示，图中展示了主数据库（master）和从数据库（slave）之间的数据复制过程。图中还标注了 8-4 页。



8-4 MySQL

MySQL 2.0 版本在 I/O 方面进行了大量的改进，MySQL 4.0 版本在 SQL 方面进行了大量的改进，MySQL 1.0 版本在 SQL 方面进行了大量的改进，MySQL 1.0 版本在 SQL 方面进行了大量的改进。

```
mysql>SHOW FULL PROCESSLIST\G;
```

```
*****1.row*****
```

```
Id:1
```

```
User:system user
```

```
Host:
```

```
db:NULL
```

```
Command:Connect
```

```
Time:6501
```

```
State:Waiting for master to send event
```

```
Info:NULL
```

```
*****2.row*****
```

```
Id:2
```

```
User:system user
```

```
Host:
```

```
db:NULL
```

```
Command:Connect
```

```
Time:0
```

```
State:Has read all relay log;waiting for the slave I/O thread to update it
```

```
Info:NULL
```

```
*****3.row*****
```

```
Id:206
```

```
User:root
```

```
Host:localhost
```

```
db:NULL
```

Command:Query

Time:0

State:NULL

Info:SHOW FULL PROCESSLIST

3 rows in set(0.00 sec)

MySQL ID 1 I/O thread
MySQL ID 2 SQL thread

replication thread

mysql>SHOW FULL PROCESSLIST\G;

.....

*****65.row*****

Id:26541

User:rep

Host:192.168.190.98:39549

db:NULL

Command:Binlog Dump

Time:6857

State:Has sent all binlog to slave;waiting for binlog to be updated

Info:NULL

.....

MySQL SHOW SLAVE STATUS SHOW MASTER STATUS

mysql>SHOW SLAVE STATUS\G;

*****1.row*****

Slave_IO_State:Waiting for master to send event

Master_Host:192.168.190.10

Master_User:rep

Master_Port:3306

Connect_Retry:60

Master_Log_File:mysql-bin.000007

Read_Master_Log_Pos:555176471

Relay_Log_File:gamedb-relay-bin.000048

Relay_Log_Pos:224355889

Relay_Master_Log_File:mysql-bin.000007

Slave_IO_Running:Yes

Slave_SQL_Running:Yes

Replicate_Do_DB:

Replicate_Ignore_DB:

Replicate_Do_Table:

Replicate_Ignore_Table:

Replicate_Wild_Do_Table:

Replicate_Wild_Ignore_Table:mysql.%,DBA.%

Last_Errno:0

Last_Error:

Skip_Counter:0

Exec_Master_Log_Pos:555176471

Relay_Log_Space:224356045

Until_Condition:None

Until_Log_File:

Until_Log_Pos:0

Master_SSL_Allowed:No

Master_SSL_CA_File:

Master_SSL_CA_Path:

Master_SSL_Cert:

Master_SSL_Cipher:

Master_SSL_Key:

Seconds_Behind_Master:0
Master_SSL_Verify_Server_Cert:No
Last_IO_Errno:0
Last_IO_Error:
Last_SQL_Errno:0
Last_SQL_Error:
1 row in set(0.00 sec)

SHOW SLAVE STATUS8-1

表 8-1 SHOW SLAVE STATUS 的主要变量

| 变 量 | 说 明 |
|-----------------------|--|
| Slave_IO_State | 显示当前 IO 线程的状态，上述状态显示的是等待主服务发送二进制日志 |
| Master_Log_File | 显示当前同步的主服务器的二进制日志，上述显示当前同步的是主服务器的 mysql-bin.000007 |
| Read_Master_Log_Pos | 显示当前同步到主服务器上二进制日志的偏移量位置，单位是字节。上述的示例显示当前同步到 mysql-bin.000007 的 555176471 偏移量位置，即已经同步了 mysql-bin.000007 这个二进制日志中 529MB (555176471/1024/1024) 的内容 |
| Relay_Master_Log_File | 当前中继日志同步的二进制日志 |
| Relay_Log_File | 显示当前写入的中继日志 |
| Relay_Log_Pos | 显示当前执行到中继日志的偏移量位置 |
| Slave_IO_Running | 从服务器中 IO 线程的运行状态，YES 表示运行正常 |
| Slave_SQL_Running | 从服务器中 SQL 线程的运行状态，YES 表示运行正常 |
| Exec_Master_Log_Pos | 表示同步到主服务器的二进制日志偏移量的位置。(Read_Master_Log_Pos - Exec_Master_Log_Pos) 可以表示当前 SQL 线程运行的延时，单位是字节。上述例子显示当前主从服务器是完全同步的 |

```

SHOW MASTER STATUS

```

```
mysql> SHOW MASTER STATUS\G;
```

*****1.row*****

File:mysql-bin.000007

Position:606181078

Binlog_Do_DB:

Binlog_Ignore_DB:

1 row in set(0.01 sec)

606181078
Read_Master_Log_PosI/O

MySQLI/OSQL
SQL
SQL

8.7.2 安全+高可用性

数据库安全+高可用性

MySQL数据库安全+高可用性

DNS Round-Robin Linux LVS

DBA DROP DATABASE DROP TABLE

point-in-time 8-5

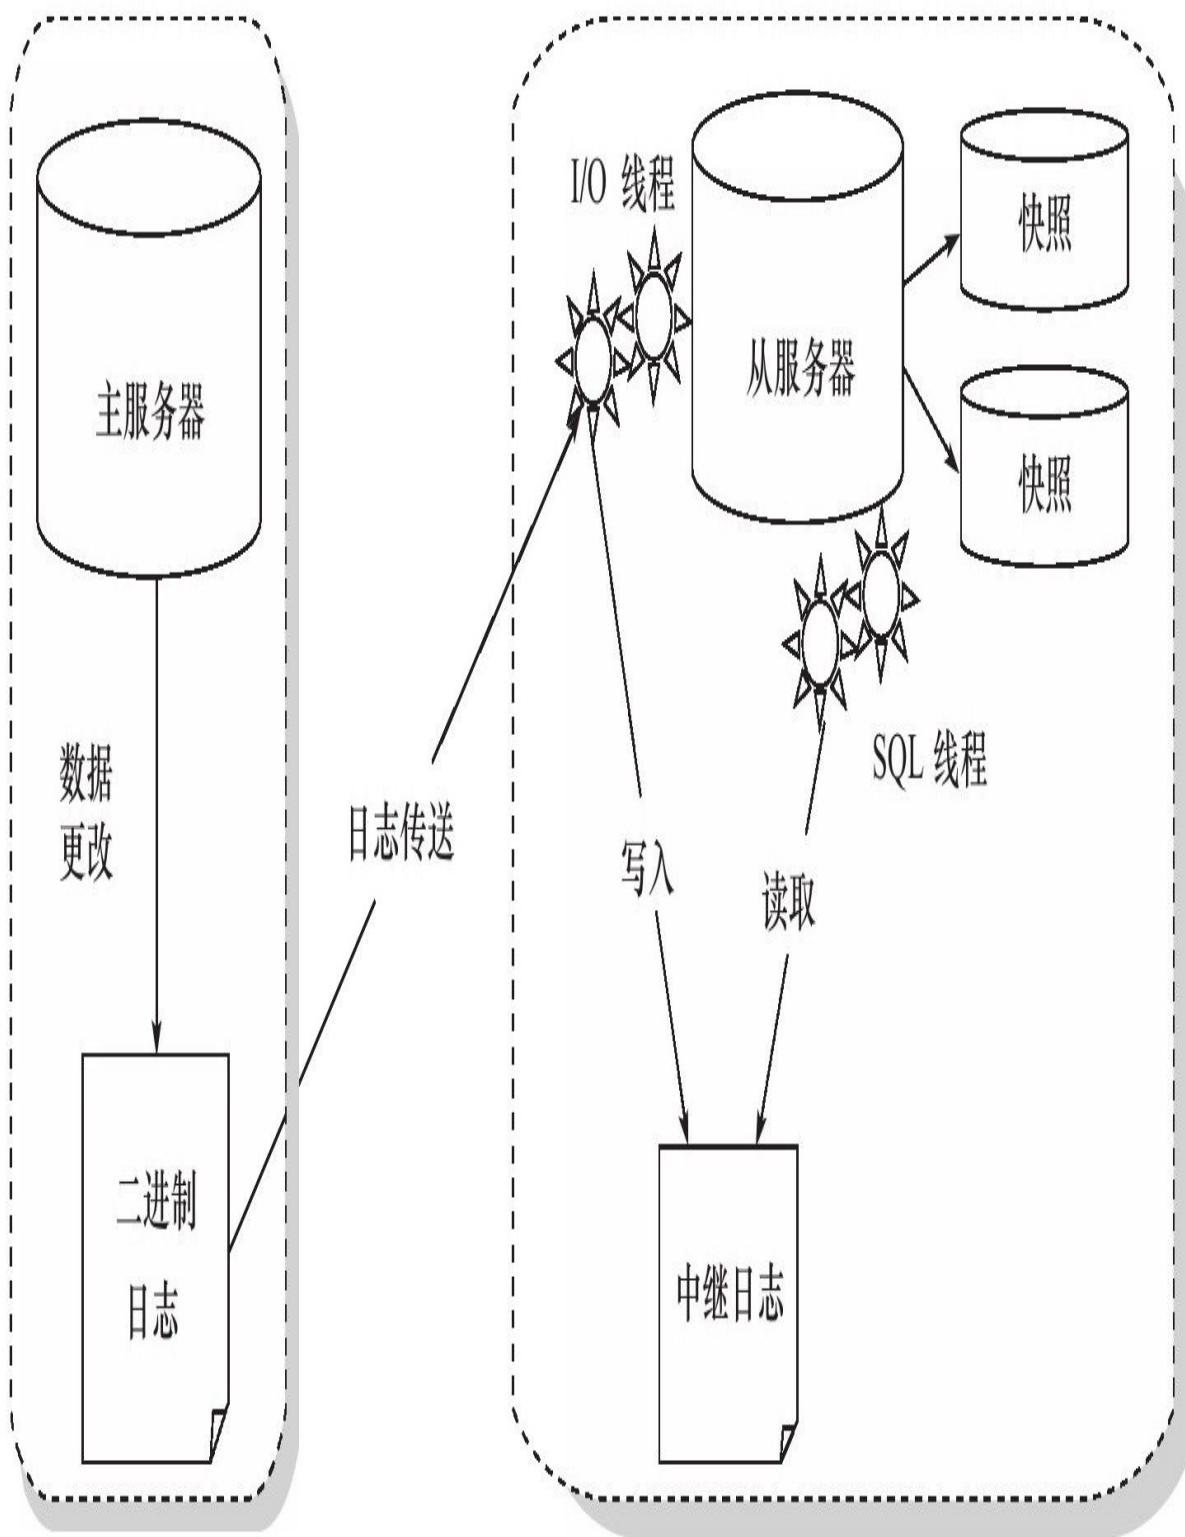


图 8-5 主从复制架构

mysql> show variables like '%read-only%';

read-only read-only

[mysqld]

read-only

mysql> show variables like '%read-only%';

mysql> INSERT INTO z SELECT 2;

ERROR 1290 (HY000): The MySQL server is running with the --read-only option so it cannot execute this statement

8.8 备份

MySQL 数据库的备份和恢复是数据库管理员的重要任务之一。MySQL 提供了多种备份和恢复工具，包括 mysqldump、mybackup 和 mysqlbackup 等。本章将介绍 MySQL 数据库的备份和恢复方法。

InnoDB 存储引擎支持事务日志，因此可以使用 mysqldump 和 xtrabackup 等工具进行备份。mysqldump 是一个命令行工具，用于导出数据库中的数据。xtrabackup 是一个第三方工具，用于备份 InnoDB 存储引擎的数据库。本章将介绍如何使用 mysqldump 和 xtrabackup 进行备份。

MySQL 数据库的恢复方法包括使用 mysqldump 和 xtrabackup 进行恢复。本章将介绍如何使用 mysqldump 和 xtrabackup 进行恢复。

9 数据库

数据库是数据库管理系统（InnoDB）的集合。数据库是数据库管理系统（InnoDB）的集合。数据库是数据库管理系统（InnoDB）的集合。

☐ 数据库CPU

☐ 数据库

☐ 数据库数据库

☐ 数据库RAID

☐ 数据库数据库

☐ 数据库数据库数据库

☐ 数据库数据库数据库

9.1 如何調度CPU

OLTP (Online Transaction Processing) 與 OLAP (Online Analytical Processing) 是兩種不同的數據處理模式。OLTP 通常用於處理大量的事務性數據，而 OLAP 則用於處理大量的分析性數據。SQL 是 OLTP 的標準語言，而 OLAP 則通常使用 SQL 的擴展語言。Blog 是 OLAP 的一種應用。

InnoDB 是 MySQL 的默認數據引擎，它支持 OLTP 和 OLAP 兩種數據處理模式。

如何調度 CPU

如何調度 CPU

如何調度 CPU

如何調度 CPU

OLTP 通常使用 CPU 的默認調度策略，而 OLAP 則通常使用 CPU 的實時調度策略。OLTP 的 IO 通常使用 CPU 的默認調度策略，而 OLAP 的 IO 則通常使用 CPU 的實時調度策略。

CPU 的默認調度策略是 64 個 CPU 核心，而 64 個 CPU 核心則通常使用 64 個 CPU 核心。Intel 和 AMD 的 CPU 通常支持 8 個 CPU 核心，而 128 個 CPU 核心則通常使用 128 個 CPU 核心。

InnoDB 的 master thread 是 InnoDB 的默認調度策略。

InnoDB 1.0 的 CPU 調度策略是 CPU 的默認調度策略。

InnoDB 1.2 的 purge master thread 是 InnoDB 的默認調度策略。InnoDB 1.1 的 CPU 調度策略是 innodb_read_io_threads 和 innodb_write_io_threads。IO 調度策略是 CPU 的默認調度策略。

MySQL数据库SQL数据库CPU数据库CPU数据库
OLTP数据库OLTP数据库CPU数据库
CPU数据库

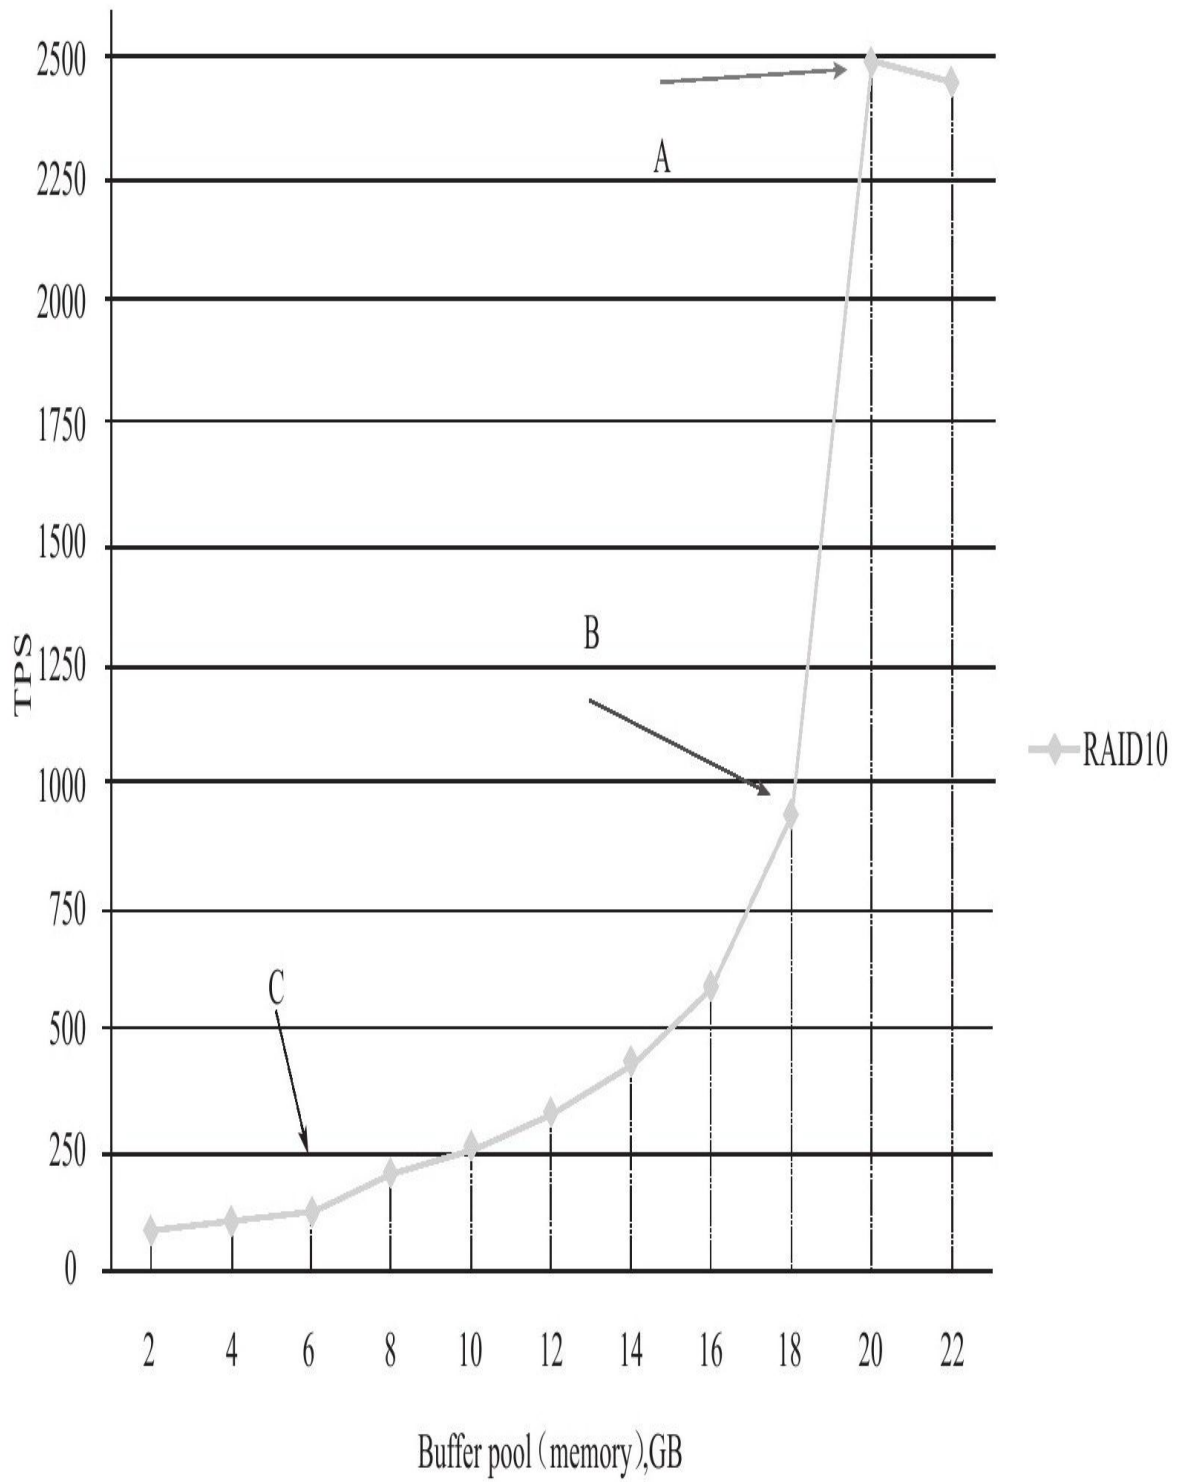
9.2 □□□□□□

```

InnoDB
InnoDB Buffer Pool
PerconaCTO Vadim
9-1

```

sysbench oltp,80mln rows (18GB data)



9-1 InnoDB

18GB 2GB 4GB 6GB 8GB 10GB 12GB 14GB 16GB 18GB 20GB 22GB sysbench TPS Transaction Per Second 20GB 22GB

“ ” 64

InnoDB 99%

```
mysql>SHOW GLOBAL STATUS LIKE'innodb%read%'\G;
```

```
*****1.row*****
```

```
Variable_name:Innodb_buffer_pool_read_ahead
```

```
Value:0
```

```
*****2.row*****
```

```
Variable_name:Innodb_buffer_pool_read_ahead_evicted
```

```
Value:0
```

```
*****3.row*****
```

```
Variable_name:Innodb_buffer_pool_read_requests
```

```
Value:167051313
```

```
*****4.row*****
```

```
Variable_name:Innodb_buffer_pool_reads
```

```
Value:129236
```

```
*****5.row*****
```

```
Variable_name:Innodb_data_pending_reads
```

```
Value:0
```

```
*****6.row*****
```

```
Variable_name:Innodb_data_read
```

```
Value:2135642112
```

*****7.row*****

Variable_name:Innodb_data_reads

Value:130309

*****8.row*****

Variable_name:Innodb_pages_read

Value:130215

*****9.row*****

Variable_name:Innodb_rows_read

Value:17651085

9 rows in set(0.00 sec)

9-1

表 9-1 当前服务器的状态参数

| 参 数 | 说 明 |
|---------------------------------------|--|
| Innodb_buffer_pool_reads | 表示从物理磁盘读取页的次数 |
| Innodb_buffer_pool_read_ahead | 预读的次数 |
| Innodb_buffer_pool_read_ahead_evicted | 预读的页，但是没有被读取就从缓冲池中被替换的页的数量，一般用来判断预读的效率 |
| Innodb_buffer_pool_read_requests | 从缓冲池中读取页的次数 |
| Innodb_data_read | 总共读入的字节数 |
| Innodb_data_reads | 发起读取请求的次数，每次读取可能需要读取多个页 |

□□□□□□□□□□□□□□□□

缓冲池命中率

$$= \frac{\text{Innodb_buffer_pool_read_requests}}{(\text{Innodb_buffer_pool_read_requests} + \text{Innodb_buffer_pool_read_ahead} + \text{Innodb_buffer_pool_reads})}$$

$$\text{平均每次读取的字节数} = \frac{\text{Innodb_data_read}}{\text{Innodb_data_reads}}$$

$$\frac{167\,051\,313}{167\,051\,313 + 129\,236 + 0} = 99.92\%$$

[illegible]

9.3 □□□□□□□□□□

9.3.1

SAS SATA 2.5 SAS

```

3ms 15 000RPM rotate per minute

```

RAID

9.3.2 闪存

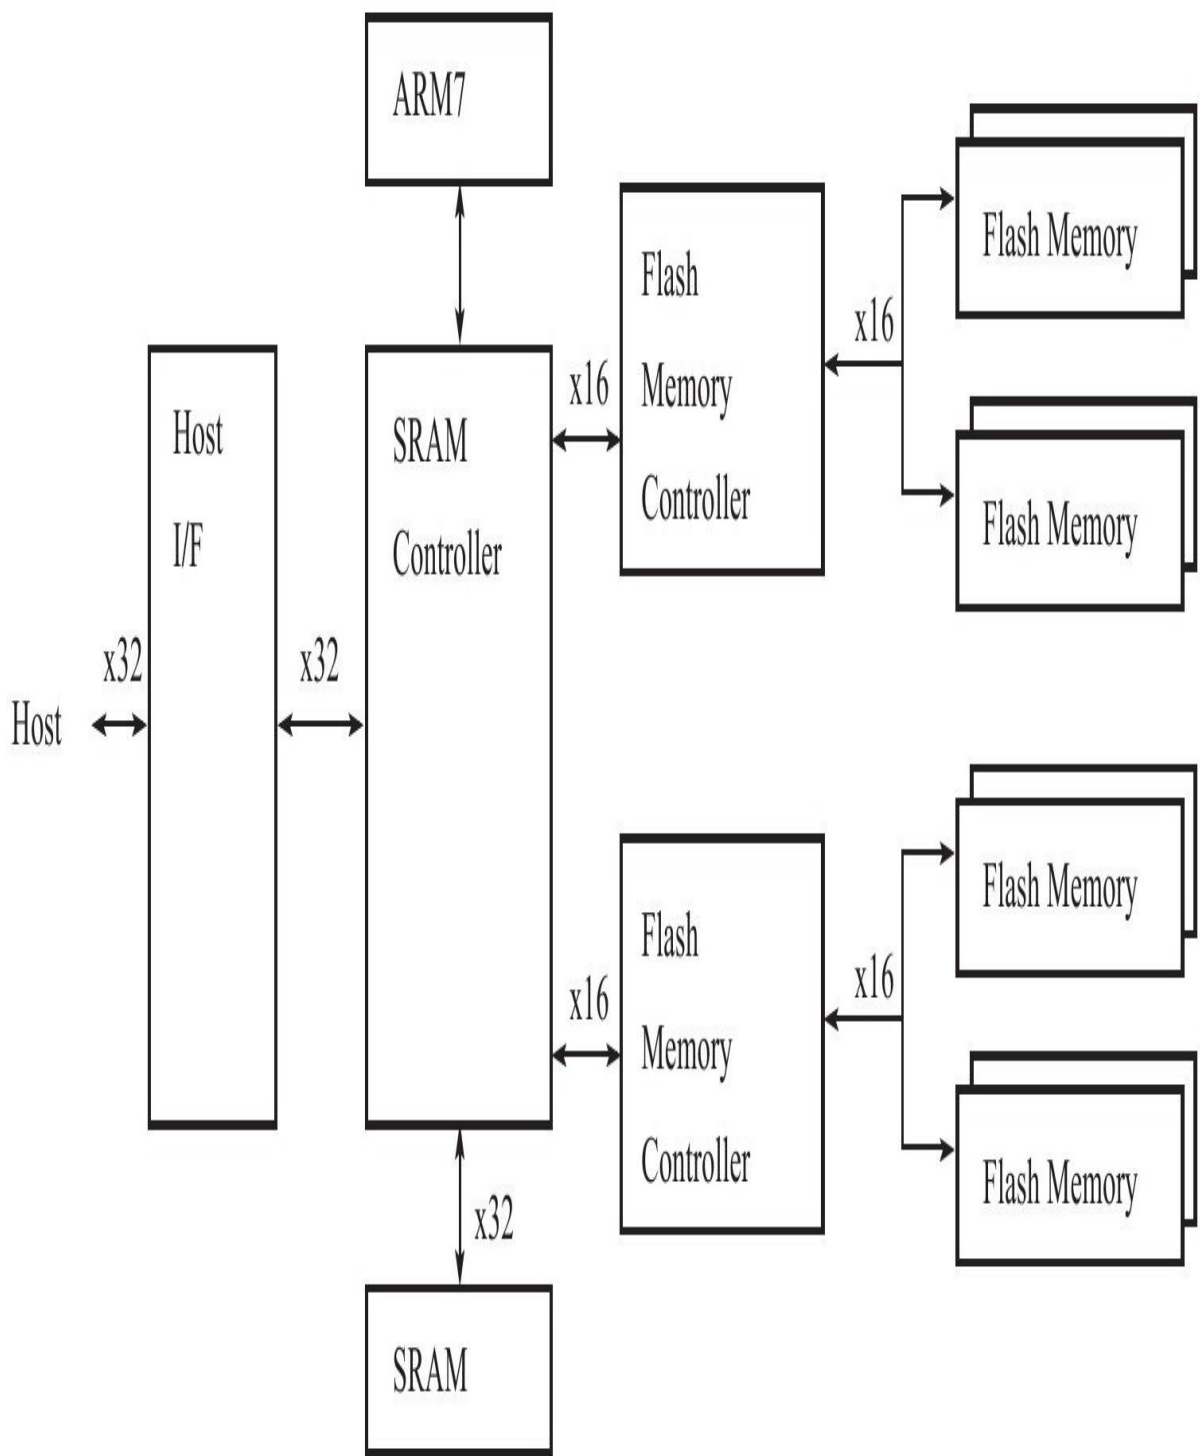
闪存（Flash Memory）是一种非易失性存储器，广泛应用于各种电子设备中。常见的闪存产品包括EMC（Enterprise Memory Cluster）、TB（Terabyte）、Oracle（Oracle Exadata）等。

闪存的主要特点是读写速度快、功耗低、体积小。它通常用于存储操作系统、应用程序和数据。闪存的容量通常以GB（Gigabyte）或TB（Terabyte）为单位。

闪存的基本单位是sector（扇区）。擦除（erase）操作通常以sector为单位进行。常见的擦除容量有128KB、256KB等。闪存的擦除速度通常比写入速度快。

闪存的寿命通常以写入次数（P/E cycles）来衡量。不同的闪存产品有不同的寿命，通常在10万次到100万次之间。

图9-2展示了闪存的基本结构。图中显示了4个闪存单元，每个单元由Intel X-25M闪存芯片组成。每个单元的数据容量为10GB。



数据库系统对I/O的访问速度要求很高，通常要求在0.1ms以内。图9-3展示了不同存储设备的随机访问时间。

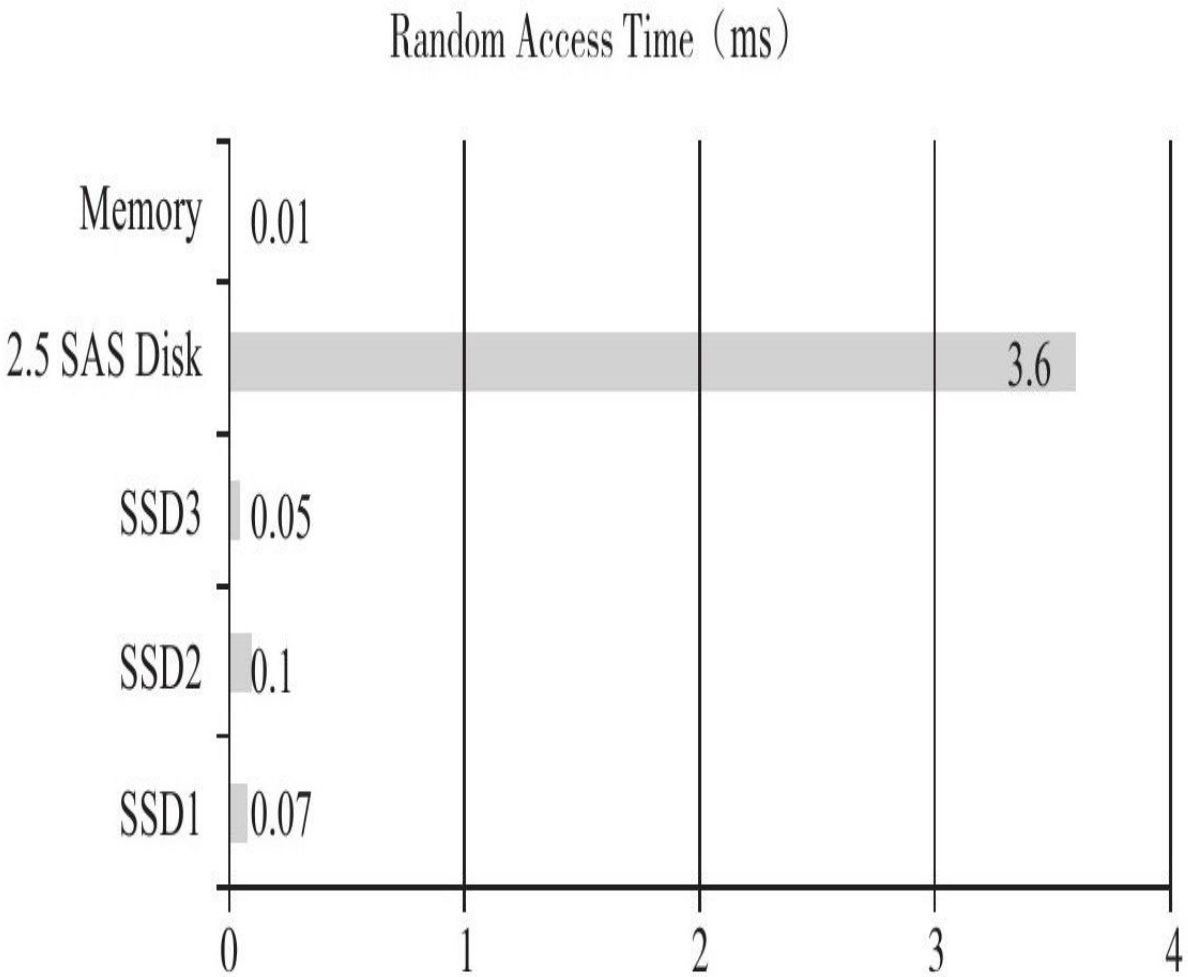


图 9-3 不同存储设备的随机访问时间

InnoDB的innodb_io_capacity参数用于控制InnoDB的IOPS。InnoDB 1.2版本中，该参数的默认值为100。InnoDB的IOPS与InnoDB的L2 Cache大小有关。

InnoDB的L2 Cache大小与InnoDB的L2 Cache大小有关。InnoDB的L2 Cache大小与InnoDB的L2 Cache大小有关。

CacheFacebook Flash
CachebcacheInnoDB

9.4 5RAID

9.4.1 RAID

RAID - Redundant Array of Independent Disks

RAID□□□□□

1 2 3 4 5 6 7 8

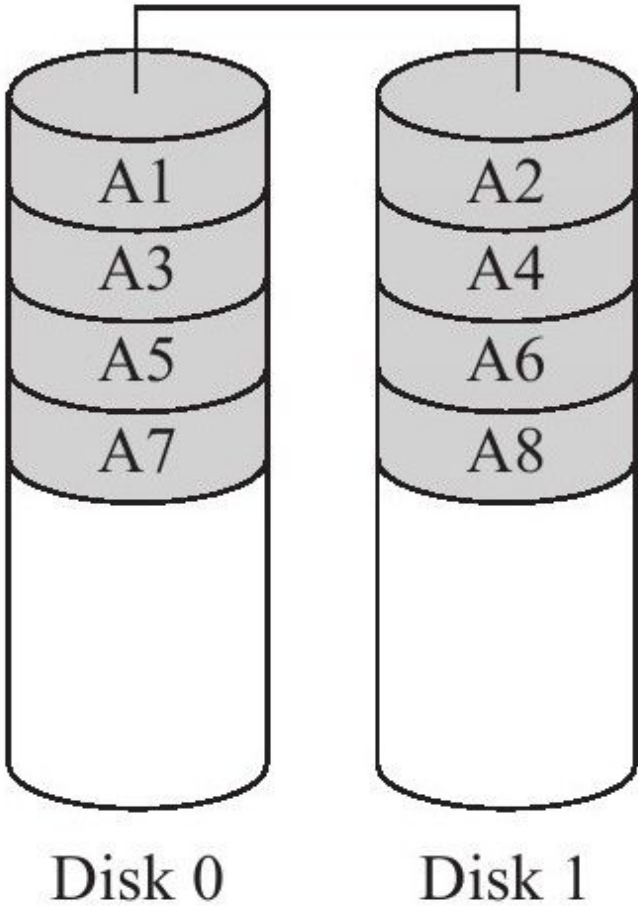
□ □ □ □ □ □

□ □ □ □ □ □ □ □

RAID 0 RAID 1 RAID 5 RAID 10 RAID 50

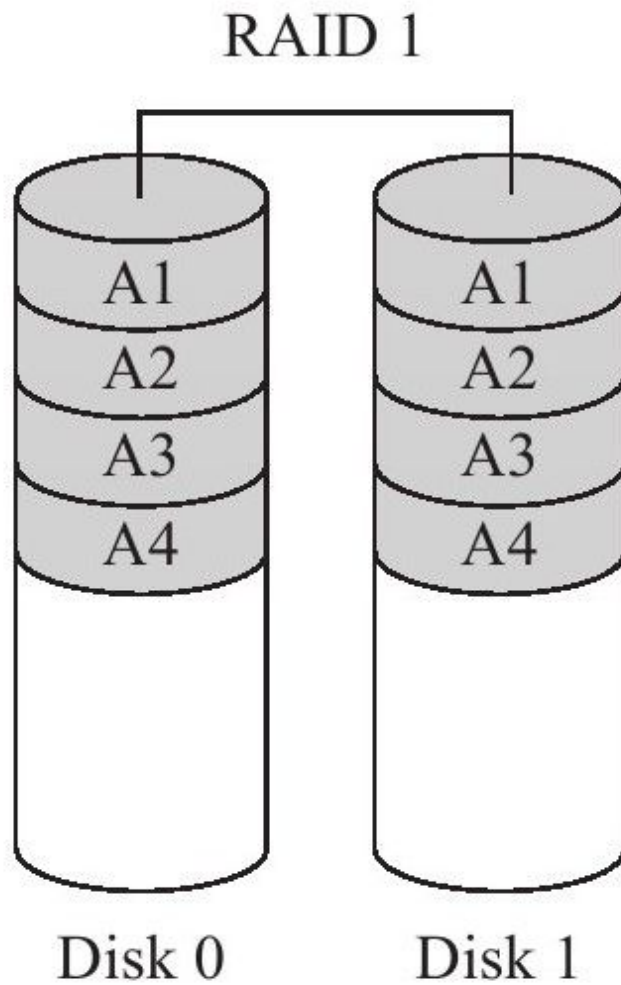
RAID 0 100MB/s 100MB/s 100MB/s 100MB/s I/O 100MB/s RAID 0 100MB/s
100MB/s 100MB/s 100MB/s 100MB/s 9-4 100MB/s 100MB/s 100MB/s
100MB/s 100MB/s 100MB/s 100MB/s RAID 0 100MB/s RAID 0 100MB/s
100MB/s 100MB/s 100MB/s 100MB/s 100MB/s 100MB/s 100MB/s 100MB/s × 100MB/s
100MB/s 100MB/s I/O 100MB/s RAID 100MB/s 100MB/s 100MB/s
100MB/s 50MB/s RAID 0 96MB/s RAID 0 130MB/s
150MB/s

RAID 0



9-4 RAID 0

[illegible]



9-5 RAID 1

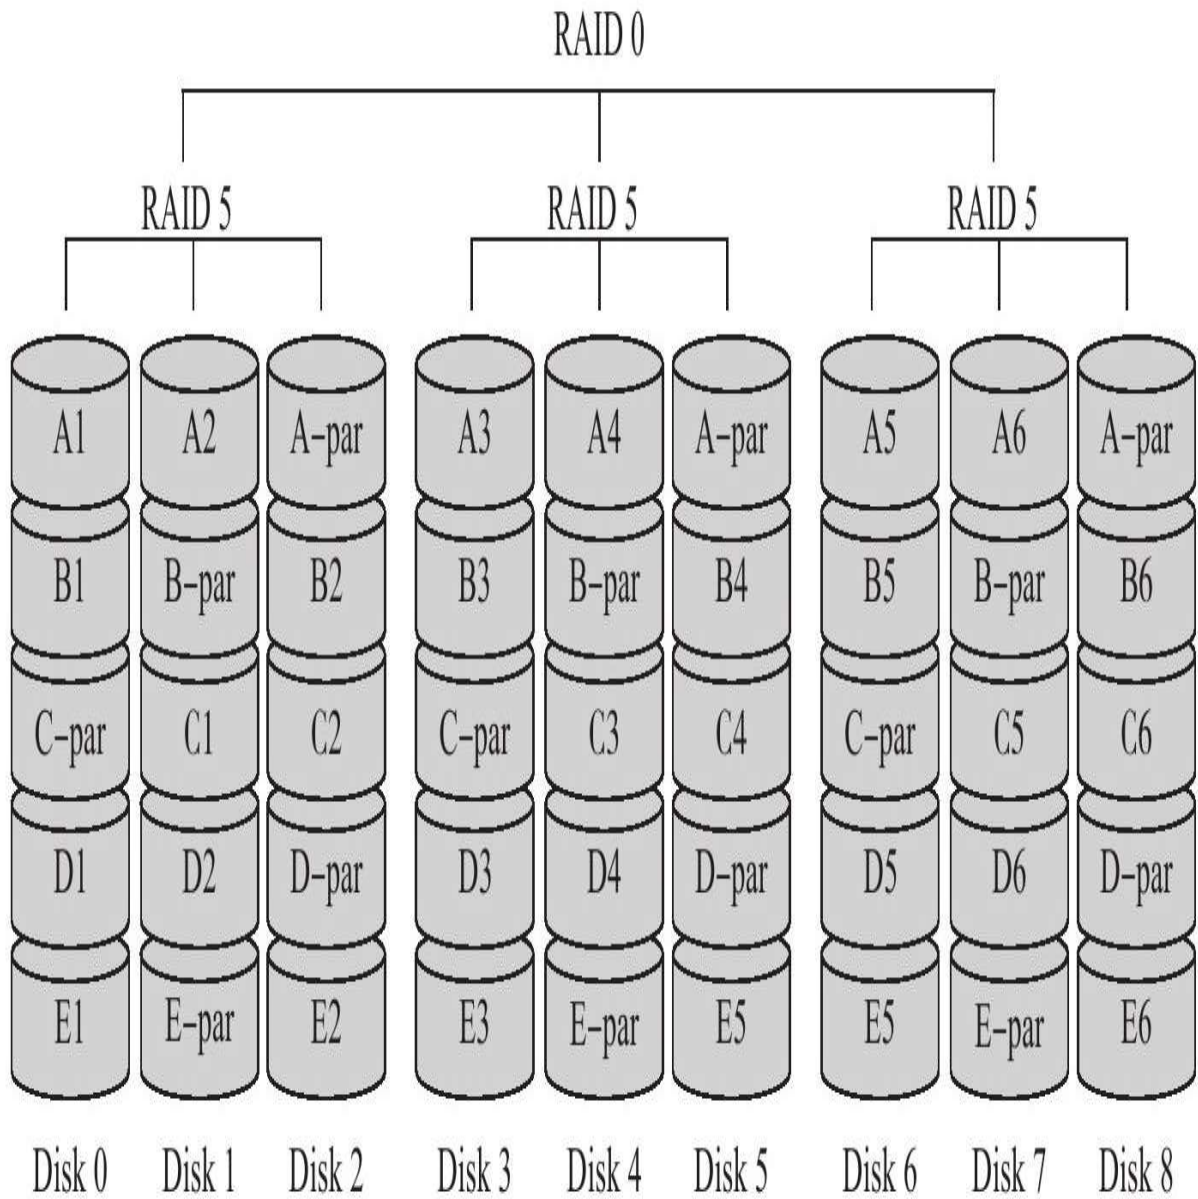
RAID 5 Disk Striping
 RAID 5 RAID 5
 RAID 5 RAID 5
 RAID 5 RAID 5
 RAID 5 RAID 5
 RAID 0 RAID 1 RAID 5
 RAID 5 RAID 0
 Write Back
 RAID 5 RAID 1 RAID 5 9-6

Diagram illustrating the layout of four disks (Disk 0, Disk 1, Disk 2, Disk 3) showing their data segments:

- Disk 0:** A1, B1, C1, Dp
- Disk 1:** A2, B2, Cp, D1
- Disk 2:** A3, Bp, C2, D2
- Disk 3:** Ap, B3, C3, D3

RAID 10 RAID 01 RAID 10 RAID 0 RAID 0
RAID 1 RAID 10 RAID 0 RAID 0
RAID 01 RAID 10 RAID 0 RAID
01 RAID 1 RAID 0 RAID
01 RAID 10 RAID 0 RAID 0
RAID 01 RAID 10 RAID 0 RAID
10 RAID 0 RAID 1 RAID 0
RAID 10 RAID 01 RAID 01 9-7

RAID 50



□ 9-8 RAID 50 □ □

□ □ □ □ □ □ □ □ RAID 10 □ □ □ □ □ □ □ □ □ □ RAID 1 □ RAID 0 □ □ □ □ □ □ □ □
 □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ strip □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □ □

IO 20% IO 90%

9.4.2 RAID Write Back

[illegible]

Write Back
RAID
RAID
BBU Battery Backup Unit
Write Back
RAID

Write Back RAID Write Through Write Through

```
Write Back RAID Write Through
Write Back RAID RAID
Write Back
Write Through
```

Write Back Write Back Write Back

□□□□20W□□□□Write Back□Write Through□□□□

```
mysql>CREATE TABLE t(a CHAR(2))Engine=InnoDB;
```

Query OK, 0 rows affected (0.00 sec)

```
mysql> DELIMITER //
```

mysql

mysql CREATE PROCEDURE p()

- □ BEGIN

```
-□DECLARE v INT;
```

```

-SET v=0;

-WHILE v<200000 DO

-INSERTINTO t VALUES('aa');

-SET v=v+1;

-END WHILE;

-END

-//

Query OK,0 rows affected(0.12 sec)

mysql>DELIMITER;

```

在表t上执行20W次插入操作，比较Write Back和Write Through的性能差异。

表 9-2 Write Back 和 Write Through 的性能对比测试结果

| RAID 卡设置 | 时 间 |
|---|-------|
| Write Back | 43 秒 |
| Write Through | 31 分钟 |
| Write Through with innodb_flush_log_at_trx_commit=0 | 68 秒 |

在表t上执行20W次插入操作，比较Write Back和Write Through的性能差异。

Write Through性能对比测试结果

innodb_flush_log_at_trx_commit=0时，性能对比测试结果

68秒

master

9.4.3 RAID

RAID BIOS RAID LSI MegaCLI

MegaCLI Windows GUI
MegaCLI Windows
MegaCLI.exe

MegaCLI RAID

```
[root@xen-server]#/opt/MegaRAID/MegaCli/MegaCli64-AdpAllInfo-a0
```

Adapter#0

=====

Versions

=====

Product Name:MegaRAID SAS 8708ELP

Serial No:P012233608

FW Package Build:9.0.1-0030

.....

HW Configuration

=====

SAS Address:500605b000d1e180

BBU:Present

Alarm:Present

NVRAM:Present

Serial Debugger:Present

Memory:Present

Flash:Present

Memory Size:256MB

TPM:Absent

Default Settings

=====

Phy Polarity:0

PhyPolaritySplit:240

Background Rate:30

Stripe Size:64kB

Flush Time:4 seconds

Write Policy:WB

Read Policy:None

Cache When BBU Bad:Disabled

Cached IO:No

SMART Mode:Mode 6

Alarm Disable:Yes

Coercion Mode:1GB

ZCR Config:Unknown

Dirty LED Shows Drive Activity:No

BIOS Continue on Error:No

Spin Down Mode:None

Allowed Device Type:SAS/SATA Mix

Allow Mix in Enclosure:Yes

Allow HDD SAS/SATA Mix in VD:Yes

Allow SSD SAS/SATA Mix in VD:No

Allow HDD/SSD Mix in VD:No

Allow SATA in Cluster:No

Max Chained Enclosures:3

Disable Ctrl-R:Yes

Enable Web BIOS:Yes

Direct PD Mapping:No

BIOS Enumerate VDs:Yes

Restore Hot Spare on Insertion:No

Expose Enclosure Devices:Yes

Maintain PD Fail History:Yes

Disable Puncturing:No

Zero Based Enclosure Enumeration:No

PreBoot CLI Enabled:Yes

LED Show Drive Activity:No

Cluster Disable:Yes

SAS Disable:No

Auto Detect BackPlane Enable:SGPIO/i2c SEP

Use FDE Only:No

Enable Led Header:No

Delay during POST:0

RAID MegaRAID SAS 8708ELP 256MB
Write Policy WB Write Back

MegaCLI

[root@xen-server]# /opt/MegaRAID/MegaCli/MegaCli64-PDList-aALL

Adapter#0

Enclosure Device ID:252

Slot Number:0

Device Id:8

Sequence Number:2

Media Error Count:0

Other Error Count:0

Predictive Failure Count:0

Last Predictive Failure Event Seq Number:0

PD Type:SAS

Raw Size:279.396 GB[0x22ecb25c Sectors]

Non Coerced Size:278.896 GB[0x22dcb25c Sectors]

Coerced Size:278.464 GB[0x22cee000 Sectors]

Firmware state:Online

SAS Address(0):0x5000c5000f363b55

SAS Address(1):0x0

Connected Port Number:0(path0)

Inquiry Data:SEAGATE ST3300655SS 00023LM5MGZZ

FDE Capable:Not Capable

FDE Enable:Disable

Secured:Unsecured

Locked:Unlocked

Foreign State:None

Device Speed:Unknown

Link Speed:Unknown

Media Type:Hard Disk Device

.....

SEAGATE ST3300655SS
<http://discountechnology.com/Seagate-ST3300655SS-SAS-Hard-Drive>
3.5 15000 Cache 16MB 3.5 4.0

Write Back

[root@xen-server]#/opt/MegaRAID/MegaCli/MegaCli64-LDGetProp-Cache-LALL-aALL

Adapter 0-VD 0(target id:0):Cache Policy:WriteBack,ReadAheadNone,Direct,No Write Cache if bad BBU

Adapter 0-VD 1(target id:1):Cache Policy:WriteBack,ReadAheadNone,Direct,No Write Cache if bad BBU

Exit Code:0x00

RAID Write Back BBU
Write Back RAID

#/opt/MegaRAID/MegaCli/MegaCli64-LDSetPropWB-LALL-aALL

#/opt/MegaRAID/MegaCli/MegaCli64-LDSetPropWT-LALL-aALL

RAID Write Back Write Through
Write Through Write Back

9.5

[illegible]

```

LinuxFreeBSDFreeBSD
MySQLFreeBSD
MySQL

```

Solaris SPARC X86
Solaris ZFS MySQL
Open Solaris

```
WindowsMySQL
WindowsMySQLLinux
WindowsLinux
```

4G 8G 4G
64 64 64
64 32 MySQL
64 64

9.6 文件系统与数据库

Windows 使用 NTFS，Solaris 使用 ZFS，Linux 使用 EXT3、EXT4、XFS、ReiserFS 等。

数据库系统通常使用文件系统来存储数据。XFS 是一种高性能的文件系统，常用于数据库。EXT3 是一种常用的文件系统，DBA 通常使用 EXT3 文件系统。

DBA 通常使用 ZFS 文件系统，LVM 通常使用 mount 文件系统。

9.7 性能测试工具

性能测试工具种类繁多，其中MySQL性能测试工具sysbench和mysql-tpcc是常用的性能测试工具。

9.7.1 sysbench

sysbench是一个跨平台的性能测试工具，支持多种数据库和操作系统。

支持的数据库：

支持的操作系统：

支持的数据库引擎：

支持的数据库表类型：

支持的数据库表引擎：

支持的数据库表引擎：

sysbench支持OLTP测试，MySQL、PostgreSQL、Oracle、Linux、Windows、Microsoft SQL Server等。

sysbench的官方网站是<http://sysbench.sourceforge.net>。Linux、RED HAT、sysbench等。

sysbench的官方网站是<http://sysbench.sourceforge.net>。

```
[root@xen-server ~]#sysbench
```

Missing required command argument.

Usage:


```
--file-total-size=SIZE total size of files to create[2G]

--file-test-mode=STRING test mode{seqwr,seqrewr,seqrd,rndrd,rndwr,rndrw}

--file-io-mode=STRING file operations mode{sync,async,fastmmap,slowmmap}[sync]

--file-extra-flags=STRING additional flags to use on opening files{sync,dsync,direct}[]

--file-fsync-freq=N do fsync()after this number of requests(0-don't use fsync())[100]

--file-fsync-all=[on|off]do fsync()after each write operation[off]

--file-fsync-end=[on|off]do fsync()at the end of test[on]

--file-fsync-mode=STRING which method to use for synchronization{fsync,fdatasync}[fsync]

--file-merged-requests=N merge at most this number of IO requests if possible(0-don't merge)[0]

--file-rw-ratio=N reads/writes ratio for combined test[1.5]
```

□□□□□□□□□□

□--file-num□□□□□□□□□□□□□□128□

□--file-block-size□□□□□□□□□□□□□□□□InnoDB□□□□□□□□
□□□□□□□□16384□□InnoDB□□□□□□□□□□□□16384□

□--file-total-size□□□□□□□□□□□□2GB□

□--file-test-mode□□□□□□□,□□seqwr□□□□□□seqrewr□□□□□□□□
seqrd□□□□□□□rndrd□□□□□□□rndwr□□□□□□□rndrw□□□□□□□□

□--file-io-mode□□□□□□□□□□□□□□□□MMAP□map□□□□□□□□
□□□□□□

□--file-extra-flags□□□□□□□□□□□□□□API□□□□□□□□

□--file-fsync-freq□□□fsync□□□□□□□fsync□□□□□□□□□□□□□□□□
□□□□□□□□□□

□--file-fsync-all□□□□□□□□□□□□□□fsync□□□□□□off□

□--file-fsync-end□□□□□□□□□□fsync□□□□□□on□

`--file-fsync-mode` 指定API `fdatasync` 还是 `fsync`

`--file-rw-ratio` 指定 `2:1`

`sysbench fileio prepare run cleanup prepare`
`run cleanup` 16 个
2GB `fileio`

```
[root@xen-server ssd]#sysbench--test=fileio--file-num=16--file-total-size=2G prepare
```

```
sysbench 0.4.10:multi-threaded system evaluation benchmark
```

```
16 files,131072Kb each,2048Mb total
```

```
Creating files for the test...
```

16 个 2GB 128MB

```
[root@xen-server ssd]#ls-lh
```

```
total 2G
```

```
-rw-----1 root root 128M Aug 12 10:42 test_file.0
-rw-----1 root root 128M Aug 12 10:42 test_file.1
-rw-----1 root root 128M Aug 12 10:42 test_file.10
-rw-----1 root root 128M Aug 12 10:42 test_file.11
-rw-----1 root root 128M Aug 12 10:42 test_file.12
-rw-----1 root root 128M Aug 12 10:42 test_file.13
-rw-----1 root root 128M Aug 12 10:42 test_file.14
-rw-----1 root root 128M Aug 12 10:42 test_file.15
-rw-----1 root root 128M Aug 12 10:42 test_file.2
-rw-----1 root root 128M Aug 12 10:42 test_file.3
-rw-----1 root root 128M Aug 12 10:42 test_file.4
-rw-----1 root root 128M Aug 12 10:42 test_file.5
-rw-----1 root root 128M Aug 12 10:42 test_file.6
```

-rw-----1 root root 128M Aug 12 10:42 test_file.7

-rw-----1 root root 128M Aug 12 10:42 test_file.8

-rw-----1 root root 128M Aug 12 10:42 test_file.9

16

[root@xen-server ssd]#sysbench--test=fileio--file-total-size=2G--file-test-mode=rndrd--max-time=180--max-requests=10000000--num-threads=16--init-rng=on--file-num=16--file-extra-flags=direct--file-fsync-freq=0--file-block-size=16384 run

100 000 000180

[root@xen-server ssd]#sysbench--test=fileio--file-total-size=2G--file-test-mode=rndrd--max-time=180--max-requests=10000000--num-threads=16--init-rng=on--file-num=16--file-extra-flags=direct--file-fsync-freq=0--file-block-size=16384 run

sysbench 0.4.10:multi-threaded system evaluation benchmark

Running the test with following options:

Number of threads:16

Initializing random number generator from timer.

Extra file open flags:16384

16 files,128Mb each

2Gb total file size

Block size 16Kb

Number of random requests for random IO:100000000

Read/Write ratio for combined random IO test:1.50

Calling fsync()at the end of test,Enabled.

Using synchronous I/O mode

Doing random read test

Threads started!

Time limit exceeded,exiting...

(last message repeated 15 times)

Done.

Operations performed:619908 Read,0 Write,0 Other=619908 Total

Read 9.459Gb Written 0b Total transferred 9.459Gb(53.81Mb/sec)

3443.85 Requests/sec executed

Test execution summary:

total time:180.0044s

total number of events:619908

total time taken by event execution:2878.0750

per-request statistics:

min:0.42ms

avg:4.64ms

max:27.30ms

approx.95 percentile:8.13ms

Threads fairness:

events(avg/stddev):38744.2500/102.69

execution time(avg/stddev):179.8797/0.00

53.81MB/s 3443.85

cleanup

[root@xen-server ssd]#sysbench--test=fileio--file-num=16--file-total-size=2G cleanup

sysbench 0.4.10:multi-threaded system evaluation benchmark

Removing test files...

#!/bin/sh

set-u

```

set-x

set-e

for size in 8G 64G;do

for mode in seqrd seqrw rndrd rndwr rndrw;do

for blksize in 4096 16384;do

sysbench--test=fileio--file-num=64--file-total-size=$size prepare

for threads in 1 4 8 16 32;do

echo"=====testing$blksize in$threads threads"

echo PARAMS$size$mode$threads$blksize sysbench-size-$size-mode-$mode-threads-$threads-blksz-$blksize

for i in 1 2 3;do

sysbench--test=fileio--file-total-size=$size--file-test-mode=$mode\

--max-time=180--max-requests=100000000--num-threads=$threads--init-rng=on\

--file-num=64--file-extra-flags=direct--file-fsync-freq=0--file-block-size=$blksize run\

|tee-a sysbench-size-$size-mode-$mode-threads-$threads-blksz-$blksize 2>>1

done

done

sysbench--test=fileio--file-total-size=$size cleanup

done

done

done

```

MySQL OLTP fileio prepare run cleanup
 prepare sbtest sbtest
 sysbench InnoDB 8000W

```

[root@xen-server]#sysbench--test=oltp--oltp-table-size=80000000--db-driver=mysql--mysql-socket=/tmp/mysql.sock--mysql-
user=root prepare

```

sysbench 0.4.10:multi-threaded system evaluation benchmark

Creating table'sbtest'...

Creating 80000000 records in table'sbtest'...

sysbench oltp

```
sysbench--test=oltp--oltp-table-size=80000000--oltp-read-only=off--init-rng=on--num-threads=16--max-requests=0--oltp-dist-type=uniform--max-time=3600--mysql-user=root--mysql-socket=/tmp/mysql.sock--db-driver=mysql run[]res
```

sysbench res[]res[]

sysbench 0.4.10: multi-threaded system evaluation benchmark

WARNING: Preparing of "BEGIN" is unsupported, using emulation

(last message repeated 15 times)

Running the test with following options:

Number of threads: 16

Initializing random number generator from timer.

Doing OLTP test.

Running mixed OLTP test

Using Uniform distribution

Using "BEGIN" for starting transactions

Using auto_inc on the id column

Threads started!

Time limit exceeded, exiting...

(last message repeated 15 times)

Done.

OLTP test statistics:

queries performed:

read: 6043324

write: 2158330

other: 863332

total: 9064986

transactions: 431666 (119.90 per sec.)

deadlocks: 0 (0.00 per sec.)

read/write requests:8201654(2278.07 per sec.)

other operations:863332(239.80 per sec.)

Test execution summary:

total time:3600.2672s

total number of events:431666

total time taken by event execution:57598.5965

per-request statistics:

min:6.84ms

avg:133.43ms

max:7155.61ms

approx.95 percentile:325.84ms

Threads fairness:

events(avg/stddev):26979.1250/64.14

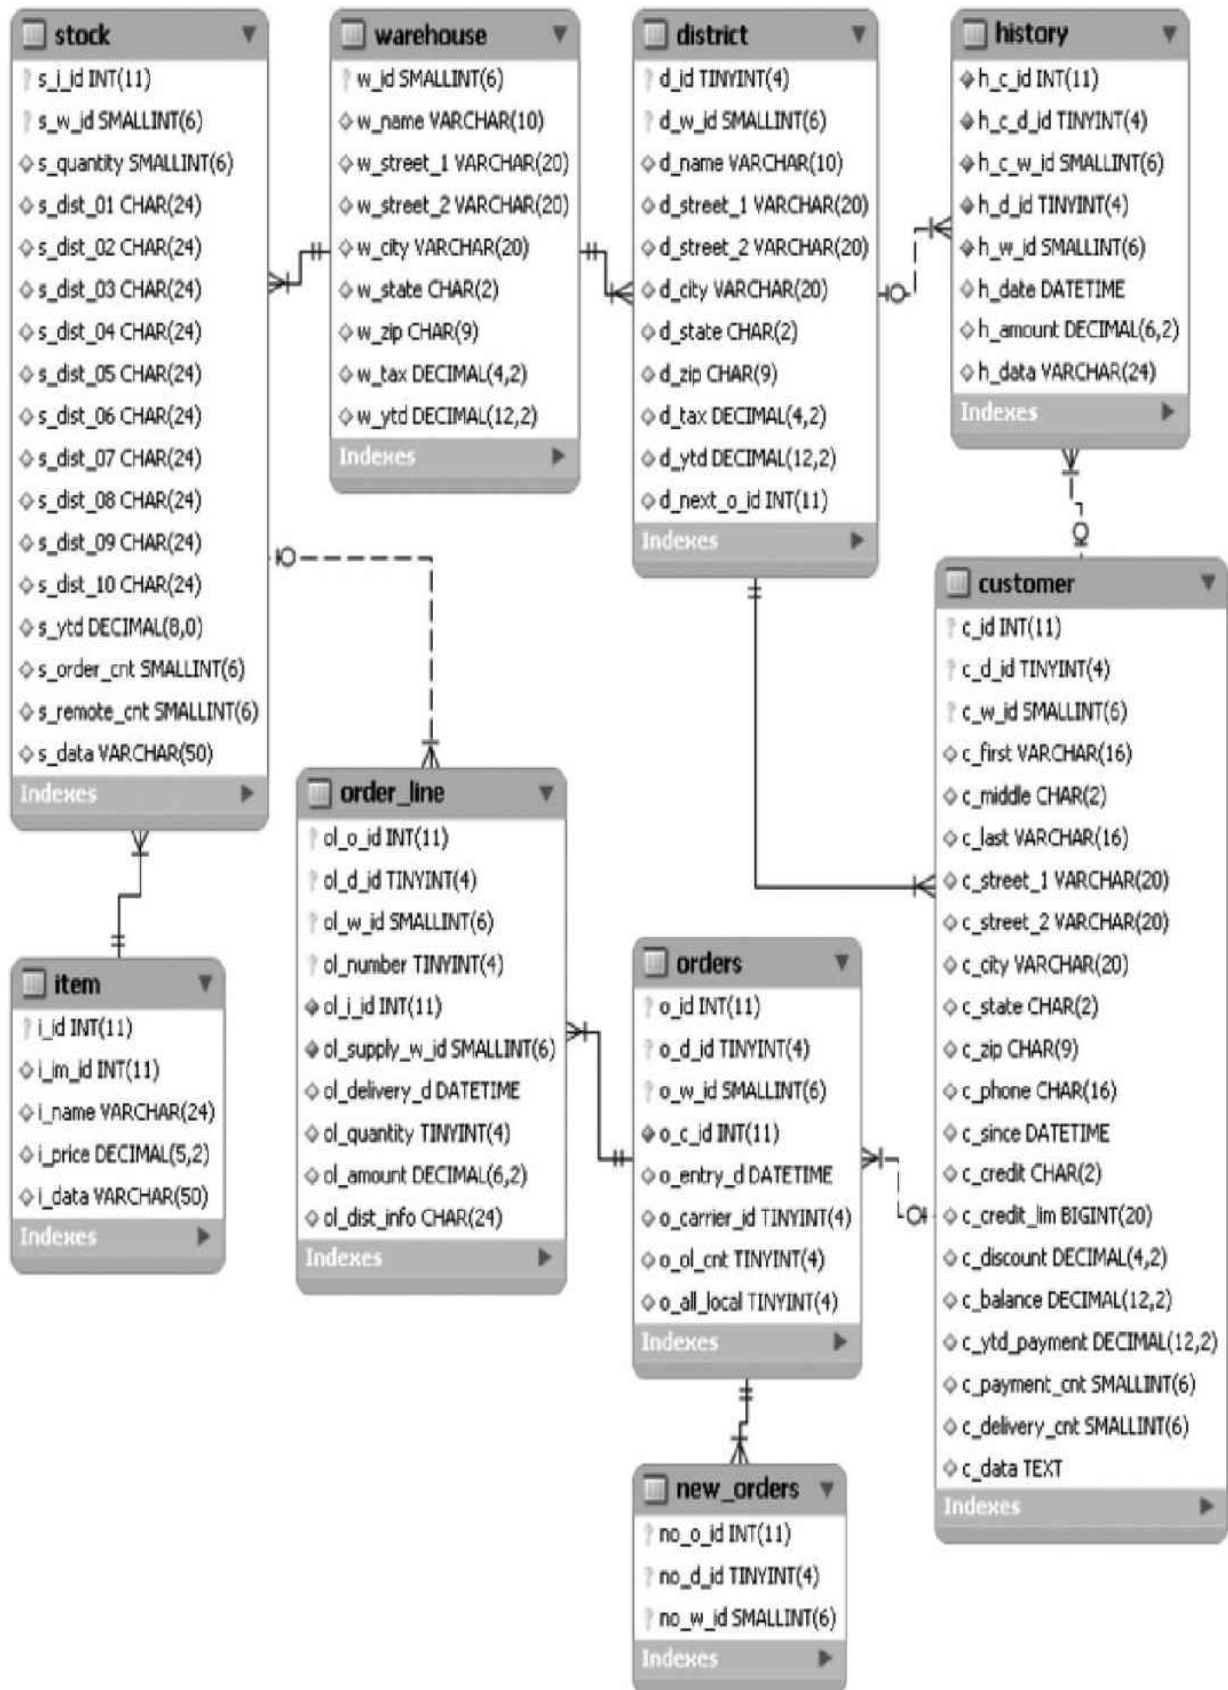
execution time(avg/stddev):3599.9123/0.06.

transactionsTPS
119.9tps sysbench OLTP
TPS sysbench

9.7.2 mysql-tpcc

TPC Transaction Processing Performance Council
TPC-C TPC OLTP TPC-C OLTP

TPC-C 3NF 10 3000 100 000 9 9-9



9-9 TPC-C

TPC-C transaction per minute (tpmC) transaction per minute (tpmC)
TPC-C

tpcc-mysql TPC-C
<https://code.launchpad.net/percona-dev/perconatools/tpcc-mysql>
tpcc-mysql Linux Windows
<http://code.google.com/p/david-mysql-tools/downloads/list> Windows tpcc-mysql

tpcc-mysql

tpcc_load 9

tpcc_start TPC-C

tpcc_load

```
[root@xen-server]#tpcc_load

*****

***##easy##TPC-C Data Loader***

*****

usage:tpcc_load[server][DB][user][pass][warehouse]

OR

tpcc_load[server][DB][user][pass][warehouse][part][min_wh][max_wh]

*[part]:1=ITEMS 2=WAREHOUSE 3=CUSTOMER 4=ORDERS
```

server MySQL IP

DB

user MySQL

passMySQL

warehouse

tpcc_load100tpcc

```
[root@xen-server tpcc-mysql]#mysql tpcc<create_table.sql
```

```
[root@xen-server tpcc-mysql]#mysql tpcc<add_fkey_idx.sql
```

```
[root@xen-server tpcc-mysql]#tpcc_load 127.0.0.1 tpcc2 root xxxxx 100
```

```
*****
```

```
***##easy##TPC-C Data Loader***
```

```
*****
```

```
<Parameters<
```

```
[server]:127.0.0.1
```

```
[DBname]:tpcc2
```

```
[user]:root
```

```
[pass]:
```

```
[warehouse]:100
```

```
TPCC Data Load Started...
```

```
Loading Item
```

```
.....5000
```

```
.....10000
```

```
.....15000
```

```
.....
```

```
...DATA LOADING COMPLETED SUCCESSFULLY.
```

tpcc_start

```
[root@xen-server]>#tpcc_start
```

```
*****
```

```
***##easy##TPC-C Load Generator***
```

usage: tpcc_start [server] [DB] [user] [pass] [warehouse] [connection] [rampup] [measure]

□□□□□□□□□□

□connection□□□□□□□□□□

□rampup□□□□□□□□□□□□□□□□□□□□□□

□measure□□□□□□□□□□

□□□tpcc_start□□16□□□□□□□□□□□□10□□□□□□□□□□20□□□□□□□□

[root@xen-server]#tpcc_start 127.0.0.1 tpcc root xxxxxx 100 16 600 1200

##easy##TPC-C Load Generator

□Parameters□

[server]:127.0.0.1

[DBname]:tpcc

[user]:root

[pass]:xxxxxx

[warehouse]:100

[connection]:16

[rampup]:600(sec.)

[measure]:1200(sec.)

.....

□□□□□□□□□□□□□□□□□□□□□□□□□□□□

RAMP-UP TIME.(1 sec.)

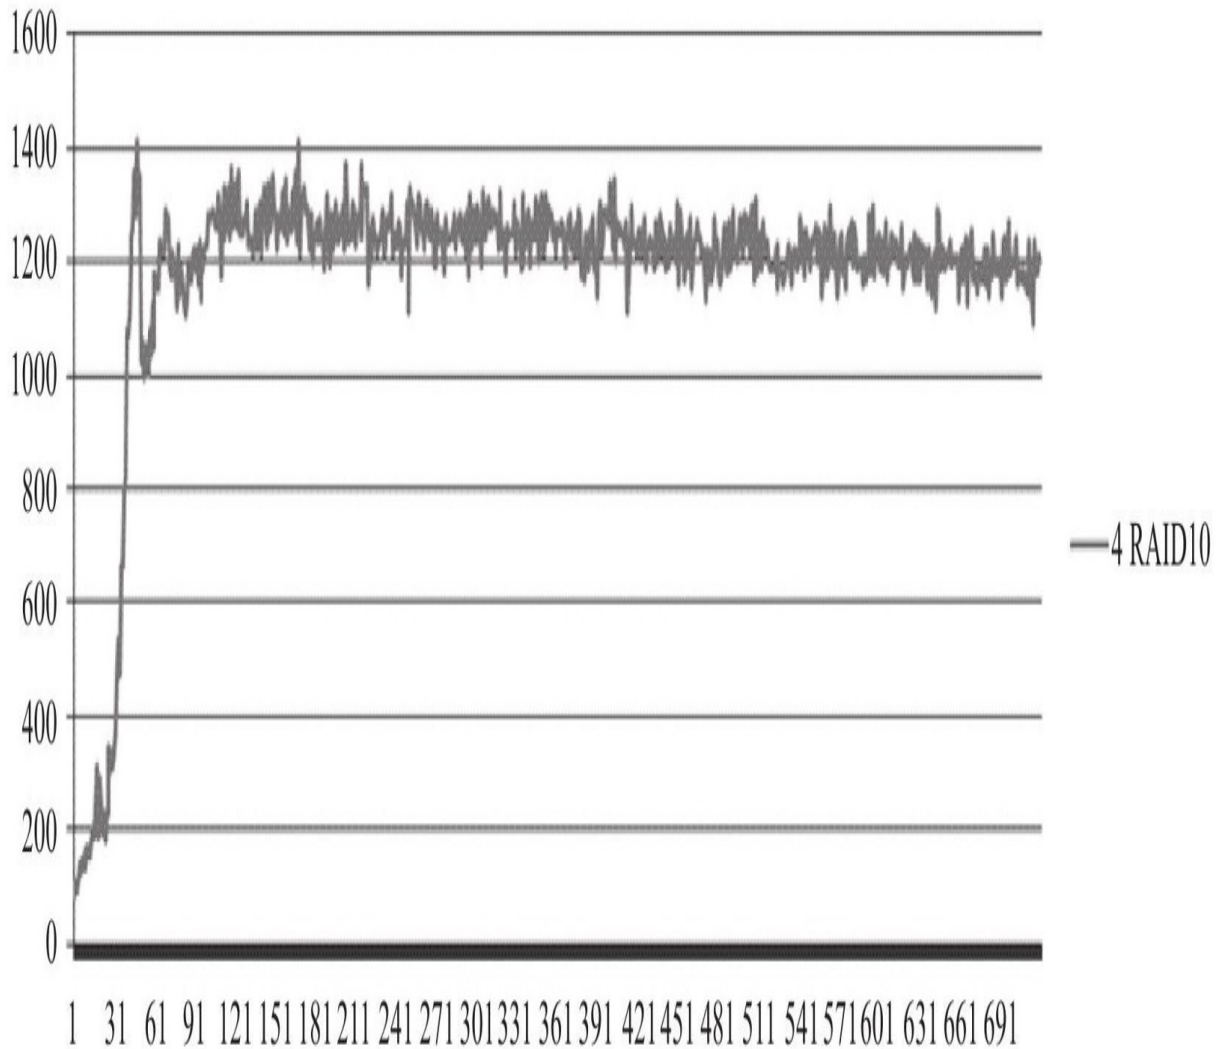
MEASURING START.

10,624(0):0.4,624(0):0.2,62(0):0.2,63(0):0.6,62(0):0.8
20,990(0):0.2,988(0):0.2,98(0):0.2,99(0):0.4,98(0):0.6
30,1435(0):0.2,1436(0):0.2,144(0):0.2,143(0):0.2,144(0):0.4
40,1736(0):0.2,1739(0):0.2,174(0):0.2,174(0):0.2,174(0):0.4
50,2041(0):0.2,2044(0):0.2,204(0):0.2,204(0):0.2,207(0):0.2
60,2195(0):0.2,2193(0):0.2,220(0):0.2,221(0):0.2,218(0):0.2
70,2332(0):0.2,2335(0):0.2,233(0):0.2,232(0):0.2,234(0):0.2
80,2408(0):0.2,2401(0):0.2,241(0):0.2,239(0):0.2,241(0):0.2
90,2473(0):0.2,2476(0):0.2,247(0):0.2,250(0):0.2,248(0):0.2
100,2350(0):0.2,2347(0):0.2,235(0):0.2,233(0):0.2,235(0):0.2

.....

TPC-C 10 TPC-C 5 New Order
Payment Order-Status Delivery Stock-Level New
Order TPC-C New Order Per 10 Second
InnoDB 10
9-10

TPCC



9-10 New Order Per 10 Second

tpcc_load tpmC New Order Per 10 Second
 New Order Per 10 Second 6 tpmC

.....
 Constraint Check(all must be[OK])

[transaction percentage]

Payment:43.48%(□=43.0%) [OK]

Order-Status:4.35%(□=4.0%) [OK]

Delivery:4.35%(□=4.0%) [OK]

Stock-Level:4.35%(□=4.0%) [OK]

[response time(at least 90%passed)]

New-Order:99.72%[OK]

Payment:99.95%[OK]

Order-Status:99.93%[OK]

Delivery:100.00%[OK]

Stock-Level:100.00%[OK]

□TpmC□

7949.942 TpmC

9.8 调优

本文档主要介绍InnoDB数据库的CPU、RAID、Linux、sysbench、tpcc-mysql、MySQL的调优方法。

10 InnoDB

InnoDB Why
What
InnoDB

10.1 InnoDB

InnoDB MySQL MySQL [\[1\]](#) MySQL 10-1

MySQL Community Server 5.1.49

Select Platform:

Source Code



Select

SuSE Linux Enterprise Server ver. 11

5.1.49

22.0M

[Download](#)

(Architecture Independent), RPM Package

(MySQL-community-5.1.49-1.sles11.src.rpm)

MD5: 714c5f8bf4b1816bb88951649d9298aa

Red Hat & Oracle Enterprise Linux 5

5.1.49

22.0M

[Download](#)

(Architecture Independent), RPM Package

(MySQL-community-5.1.49-1.rhel5.src.rpm)

MD5: 8c386345d6374be174033f3a17d49a0b

SuSE Linux Enterprise Server 10 (Architecture Independent), RPM Package

5.1.49

22.0M

[Download](#)

(MySQL-community-5.1.49-1.sles10.src.rpm)

MD5: ca2ed7f15fae60331f40b0083847fe39

Generic Linux (glibc 2.3) (Architecture Independent), RPM Package

5.1.49

22.0M

[Download](#)

(MySQL-5.1.49-1.glibc23.src.rpm)

MD5: 40cb7b0399b3a174e9b9e06281f1036f

SuSE Linux Enterprise Server 9 (Architecture Independent), RPM Package

5.1.49

22.0M

[Download](#)

(MySQL-community-5.1.49-1.sles9.src.rpm)

MD5: 9e171c70c47d6792b0856021bbd12353

Red Hat & Oracle Enterprise Linux 4

5.1.49

22.0M

[Download](#)

(Architecture Independent), RPM Package

(MySQL-community-5.1.49-1.rhel4.src.rpm)

MD5: 897d2e9ffbe9109072fe22bcafa91788

Red Hat Enterprise Linux 3 (Architecture Independent), RPM Package

5.1.49

22.0M

[Download](#)

(MySQL-community-5.1.49-1.rhel3.src.rpm)

MD5: bc855486bd2b4d5d029d7b6bba4d4361

Generic Linux (Architecture Independent),

5.1.49

22.6M

[Download](#)

10-1 MySQL

Generic Linux
MySQL Download GA
MySQL MySQL 5.5.5 milestone GA
www dev
http://dev.mysql.com/downloads/mysql MySQL
10-2

[Generally Available \(GA\) Releases](#)[Development Releases](#)

MySQL Community Server 5.5.5 m3

Select Platform:

Source Code



Select

**Linux - Generic 2.6 (Architecture Independent),
RPM Package**

(MySQL-5.5.5_m3-1.linux2.6.src.rpm)

5.5.5

21.0M

[Download](#)

MD5: ca368fb09817420b2a0f3d325f2acc34

**SuSE Linux Enterprise Server ver. 11
(Architecture Independent), RPM Package**

(MySQL-5.5.5_m3-1.sles11.src.rpm)

5.5.5

21.0M

[Download](#)

MD5: e8a13c562e1846666367d211575f2f29

**Red Hat & Oracle Enterprise Linux 5
(Architecture Independent), RPM Package**

(MySQL-5.5.5_m3-1.rhel5.src.rpm)

5.5.5

21.0M

[Download](#)

MD5: dd01ca7e34be238d62eef7c1e3d4c3fc

**SuSE Linux Enterprise Server 10 (Architecture
Independent), RPM Package**

(MySQL-5.5.5_m3-1.sles10.src.rpm)

5.5.5

21.0M

[Download](#)

MD5: d5b5e453f40e4ecb35c2731f4ca9ab2b

**Red Hat & Oracle Enterprise Linux 4
(Architecture Independent), RPM Package**

(MySQL-5.5.5_m3-1.rhel4.src.rpm)

5.5.5

21.0M

[Download](#)

MD5: 019c635b231625d206284e565079dccc

**Generic Linux (Architecture Independent),
Compressed TAR Archive**

(mysql-5.5.5-m3.tar.gz)

5.5.5

21.8M











[Download](#)

MD5: ad27f6561d6010c9346ffeca6de403fa

10-2 MySQL

```
"Download"mysql.com  
MySQL  
"Asia"10-3
```






Asia

| | |
|---|----------|
|  sPD Hosting, Israel | HTTP |
|  JAIST, Japan | HTTP FTP |
|  Internet Initiative Japan Inc., Japan | HTTP FTP |
|  Lahore University of Management Sciences, Pakistan | HTTP FTP |
|  Kyung Hee University Linux User Group, Republic of Korea | HTTP FTP |
|  ezNetworking Solutions Pte. Ltd., Singapore | HTTP FTP |
|  mirror.tw (Taiwan Mirror), Taiwan | HTTP FTP |
|  Providence University, Taiwan | HTTP FTP |
|  National Taiwan University, Taiwan | HTTP FTP |
|  National Sun Yat-Sen University, Taiwan | HTTP FTP |
|  Computer Center, Shu-Te University / Kaohsiung, Taiwan | HTTP FTP |

□ 10-3 MySQL□□□□□□

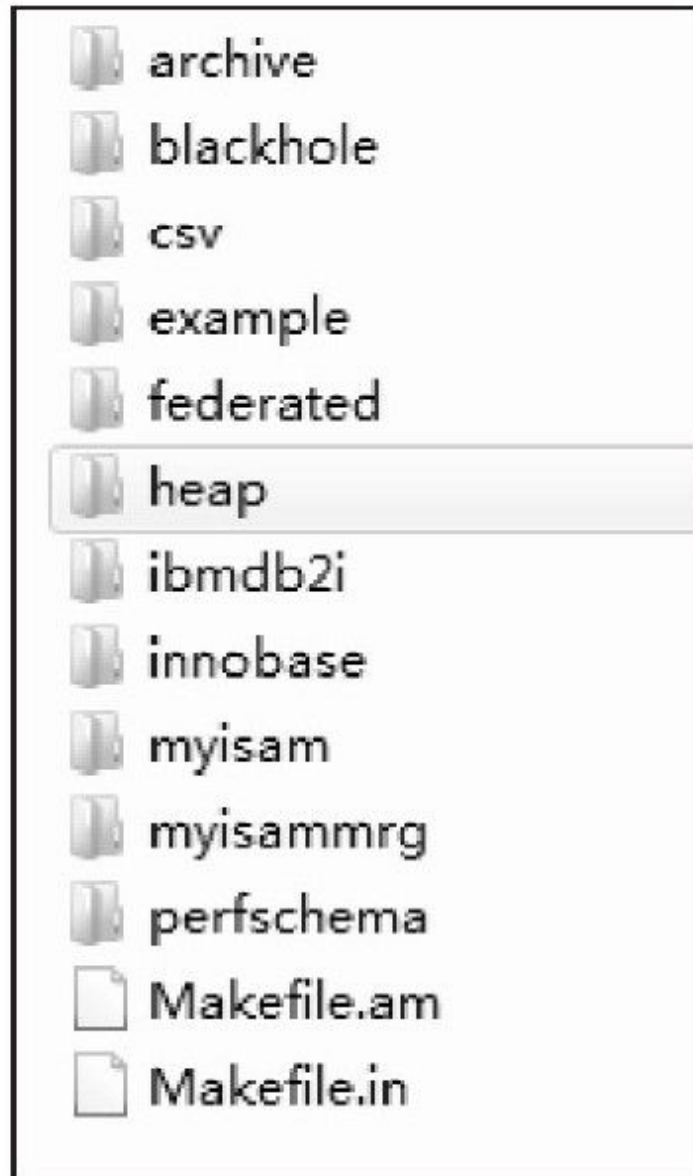
```
tar.gzLinuxtarWindowsWinRAR
MySQL10-4
```



-  BUILD
-  client
-  cmake
-  CMakeFiles
-  cmd-line-utils
-  config
-  debug
-  Docs
-  extra
-  include
-  libmysql
-  libmysql_r
-  libmysqld
-  libservices
-  man
-  mysql-test
-  mysys
-  netware
-  packaging
-  plugin
-  pstack
-  regex
-  scripts
-  sql
-  sql-bench
-  sql-common
-  storage
-  strings
-  support-files
-  tests
-  unittest
-  vio
-  win
-  zlib

□ 10-4 MySQL□□□□□□

□□□□□□□□□□storage□□□□□□□□□□□□10-5□□□



□ 10-5 □□□□□□□□

□□□□□□□□□□□□□□□□□□□□□□□□□□□□archive□
blackhole□csv□fedorated□heap□ibmdb2i□myisam□innobase□
□MySQL 5.5□□□□□InnoDB Plugin□□□□□□InnoDB□□□□□□□□
MySQL 5.1□□□□□□□□□□□□□□InnoDB□□□□□□□□□□10-6□□□



□ 10-6 MySQL 5.1□□□□□□□□

□□□□innobase□innodb_plugin□□□□□innobase□□□□□InnoDB
□□□□□□□innodb_plugin□□□□InnoDB Plugin□□□□□□□□□□
InnoDB Plugin□□□□□□MySQL□□□□□□□□□□innobase□□□□□
innodb_plugin□□□□□□innobase□

[1] <http://www.mysql.com/downloads/mysql/>

10.2 InnoDB

InnoDB 10-7
InnoDB

| | |
|---|------------|
|  | btr |
|  | buf |
|  | data |
|  | dict |
|  | dyn |
|  | eval |
|  | fil |
|  | fsp |
|  | fut |
|  | ha |
|  | handler |
|  | ibuf |
|  | include |
|  | lock |
|  | log |
|  | mach |
|  | mem |
|  | mtr |
|  | mysql-test |
|  | os |
|  | page |
|  | pars |
|  | que |
|  | read |
|  | rem |
|  | row |
|  | srv |
|  | sync |
|  | thr |
|  | trx |
|  | usr |
|  | ut |

□ 10-7 InnoDB□□□□□□□□□□

□btr□B+□□□□□

□buf□□□□□□□□□LRU□□□Flush□□□□□

□dict□InnoDB□□□□□□□□□□□□

□dyn□InnoDB□□□□□□□□□□□

□fil□InnoDB□□□□□□□□□□□□□□□□

□fsp□□□□□file space□□□InnoDB□□□□□□□□□□□□□□□□

□ha□□□□□□□□□

□handler□□□□MySQL□handler□□□□□□□□□□□

□ibuf□□□□□□□□□

□include□InnoDB□□□□□.h□.ic□□□□□□□□□□□□□

□lock□InnoDB□□□□□□□□□S□□X□□□□□□□□□□□□

□log□□□□□□□□□□□□□□□□□□□□□□□□□□□□

□mem□□□□□□□□□□□□□□□□□□□□□

□mtr□□□□□□□□□

□os□□□□□□□□□□□□□□□

□page□□□□□□□

□row□□□□□□□□□□□□□□□

□srv□□□InnoDB□□□□□□□□□□

□sync□InnoDB□□□□□□□□□Mutex□□□□□

☐thrInnoDB

☐trx

☐ut

10.3 MySQL 5.1 安装 InnoDB

10.3.1 Windows

Windows 安装 Visual Studio 2003/2005/2008
MySQL 安装

CMake <http://www.cmake.org>

bison

<http://gnuwin32.sourceforge.net/packages/bison.htm>

configure.js

```
C:\workdir\win\configure.js options
```

options

WITH_INNOBASE_STORAGE_ENGINE InnoDB

WITH_PARTITION_STORAGE_ENGINE

WITH_ARCHIVE_STORAGE_ENGINE Archive

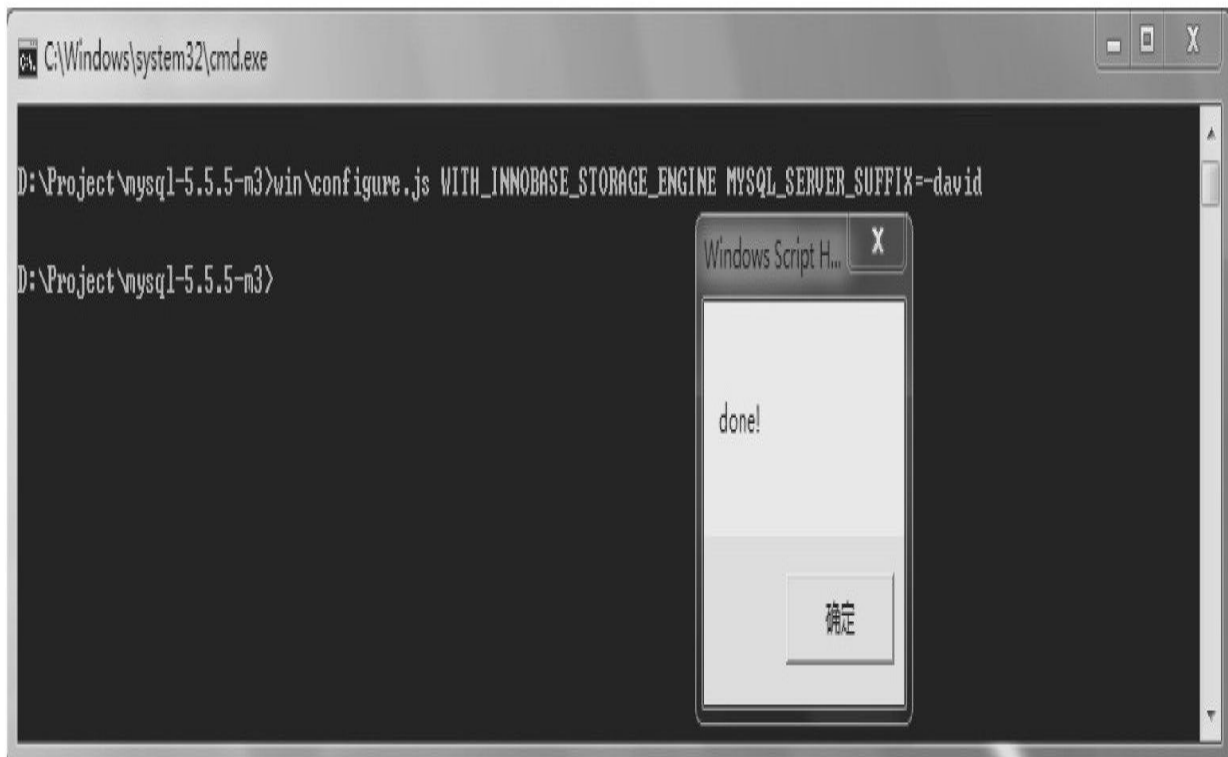
WITH_BLACKHOLE_STORAGE_ENGINE Blackhole

WITH_EXAMPLE_STORAGE_ENGINE Example
安装 InnoDB 存储引擎

WITH_FEDERATED_STORAGE_ENGINE Federated

WITH_NDBCLUSTER_STORAGE_ENGINE NDB Cluster

InnoDB 安装 10-8



□ 10-8 configure.js□□

□□□□□□□□□□Visual Studio 2005□□Visual Studio 2008□win□□
 □□□build-vs8.bat□□□□□Visual Studio□□□□□build-vs8.bat□□
 Visual Studio 2005□build-vs8_x64.bat□□□□□64□□MySQL□□□□
 □□□□□□□32□□□□□□□□Visual Studio 2008□□□□□□□□□□□□□□

D:\Project\mysql-5.5.5-m3\win\build-vs9.bat

--Check for working C compiler:C:/Program Files/Microsoft Visual Studio 9.0/VC/bin/cl.exe

--Check for working C compiler:C:/Program Files/Microsoft Visual Studio 9.0/VC/bin/cl.exe--works

--Detecting C compiler ABI info

--Detecting C compiler ABI info-done

--Check for working CXX compiler:C:/Program Files/Microsoft Visual Studio 9.0/VC/bin/cl.exe

--Check for working CXX compiler:C:/Program Files/Microsoft Visual Studio 9.0/VC/bin/cl.exe--works

--Detecting CXX compiler ABI info

--Detecting CXX compiler ABI info-done

--Check size of void*

--Check size of void*-done

sizeof_voidp=4

--Looking for include files HAVE_CXXABI_H

--Looking for include files HAVE_CXXABI_H-not found.

--Looking for include files HAVE_NDIR_H

--Looking for include files HAVE_NDIR_H-not found.

--Looking for include files HAVE_SYS_NDIR_H

--Looking for include files HAVE_SYS_NDIR_H-not found.

--Looking for include files HAVE_ASM_TERMBITS_H

--Looking for include files HAVE_ASM_TERMBITS_H-not found.

--Looking for include files HAVE_TERMBITS_H

--Looking for include files HAVE_TERMBITS_H-not found.

--Looking for include files HAVE_VIS_H

--Looking for include files HAVE_VIS_H-not found.

--Looking for include files HAVE_WCHAR_H

--Looking for include files HAVE_WCHAR_H-found

--Looking for include files HAVE_WCTYPE_H

--Looking for include files HAVE_WCTYPE_H-found

--Looking for include files HAVE_XFS_XFS_H

--Looking for include files HAVE_XFS_XFS_H-not found.

--Looking for include files CMAKE_HAVE_PTHREAD_H

--Looking for include files CMAKE_HAVE_PTHREAD_H-not found.

--Found Threads:TRUE

--Looking for pthread_rwlockattr_setkind_np

--Looking for pthread_rwlockattr_setkind_np-not found

--Performing Test HAVE_SOCKADDR_IN_SIN_LEN

--Performing Test HAVE_SOCKADDR_IN_SIN_LEN-Failed

--Performing Test HAVE_SOCKADDR_IN6_SIN6_LEN

--Performing Test HAVE_SOCKADDR_IN6_SIN6_LEN-Failed

--Cannot find wix 3,installer project will not be generated

--Configuring done

--Generating done

--Build files have been written to:D:/Project/mysql-5.5.5-m3

MySQL.sln
MySQL

Visual Studio
InnoDB master thread



□ 10-9 □□master thread

10.3.2 Linux

LinuxEclipseUNIXSolarisFreeBSD
MACEclipse

<http://www.eclipse.org/downloads/>Eclipse IDE for
C/C++DevelopersMySQL
/root/workspace/mysql-5.5.5-m3Make
EclipseMake

```
[root@xen-server mysql-5.5.5-m3]#BUILD/compile-amd64-debug-max-no-ndb-c
```

BUILDcompile64Linux
Debugcompile-amd64-debug-max-no-
ndb-cMake

EclipseC++10-10



□ 10-10 □□□C++□□

□□□□□□□□□□□□mysql_5_5_5□□□□□□□□□□10-11□□□

C++ Project

C++ Project

Create C++ project of selected type

Project name:

mysql_5_5_5

☒ Use default location

Location:

/root/workspace/mysql_5_5_5

Browse...

Project type:

Toolchains:

▼ Executable

Empty Project

● Hello World C++ Project

▶ Shared Library

▶ Static Library

▶ Makefile project

Linux GCC

☒ Show project types and toolchains only if they are supported on the platform

?

< Back

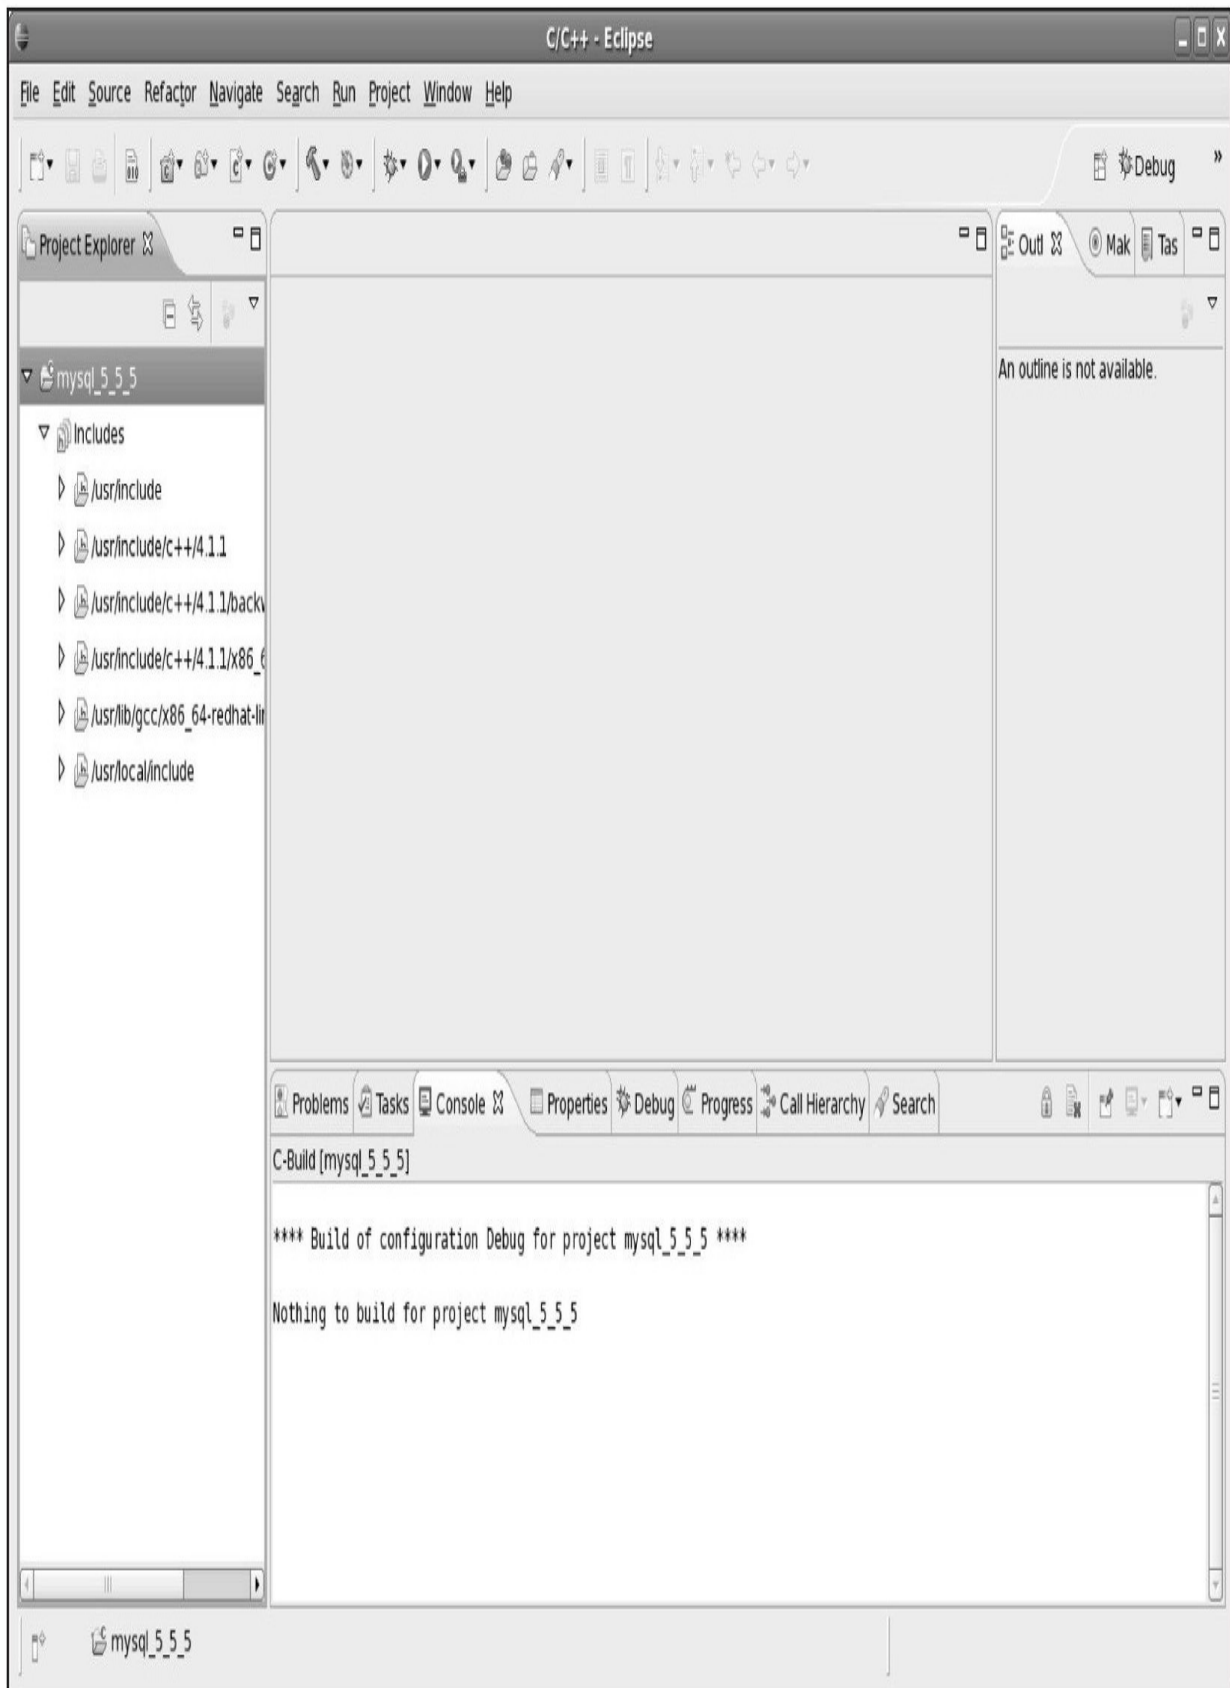
Next >

Cancel

Finish

□ 10-11 □□□□

□□Finish□□□□□□□□□□□□□□□□10-12□□□



□ 10-12 □□□C++□□

□□□□□□Project Explorer□□□□□mysql_5_5_5□□□□□□□□□□
□/root/workspace/mysql-5.5.5-m3□□□□□□□10-13□□□

New Folder

Folder

Create a new folder resource.



Enter or select the parent folder:

mysql_5_5_5



 mysql_5_5_5

Folder name:

mysql-5.5.5-m3

<< Advanced

☒ Link to folder in the file system

/root/workspace/mysql-5.5.5-m3

Browse...

Variables...

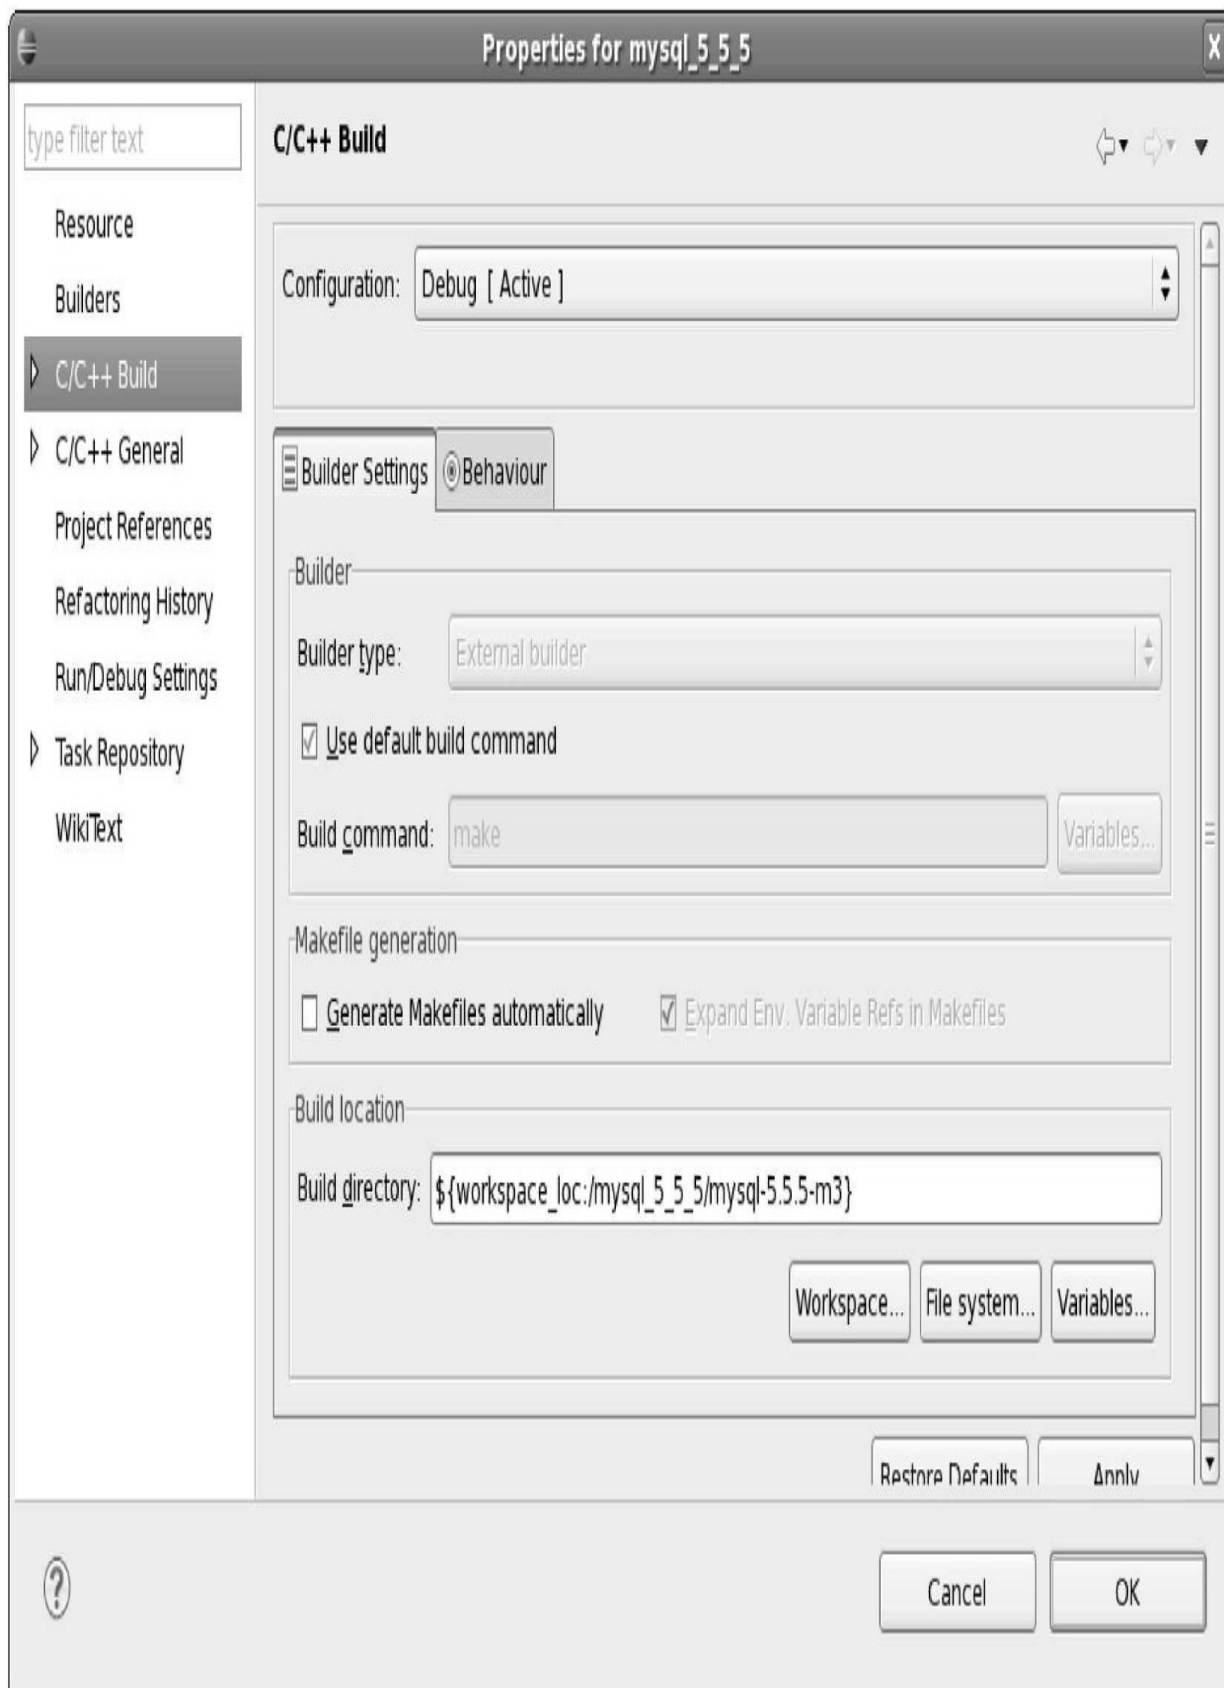


Cancel

Finish

10-13

mysql_5_5_5 C/C++ Build
Build directory 10-14



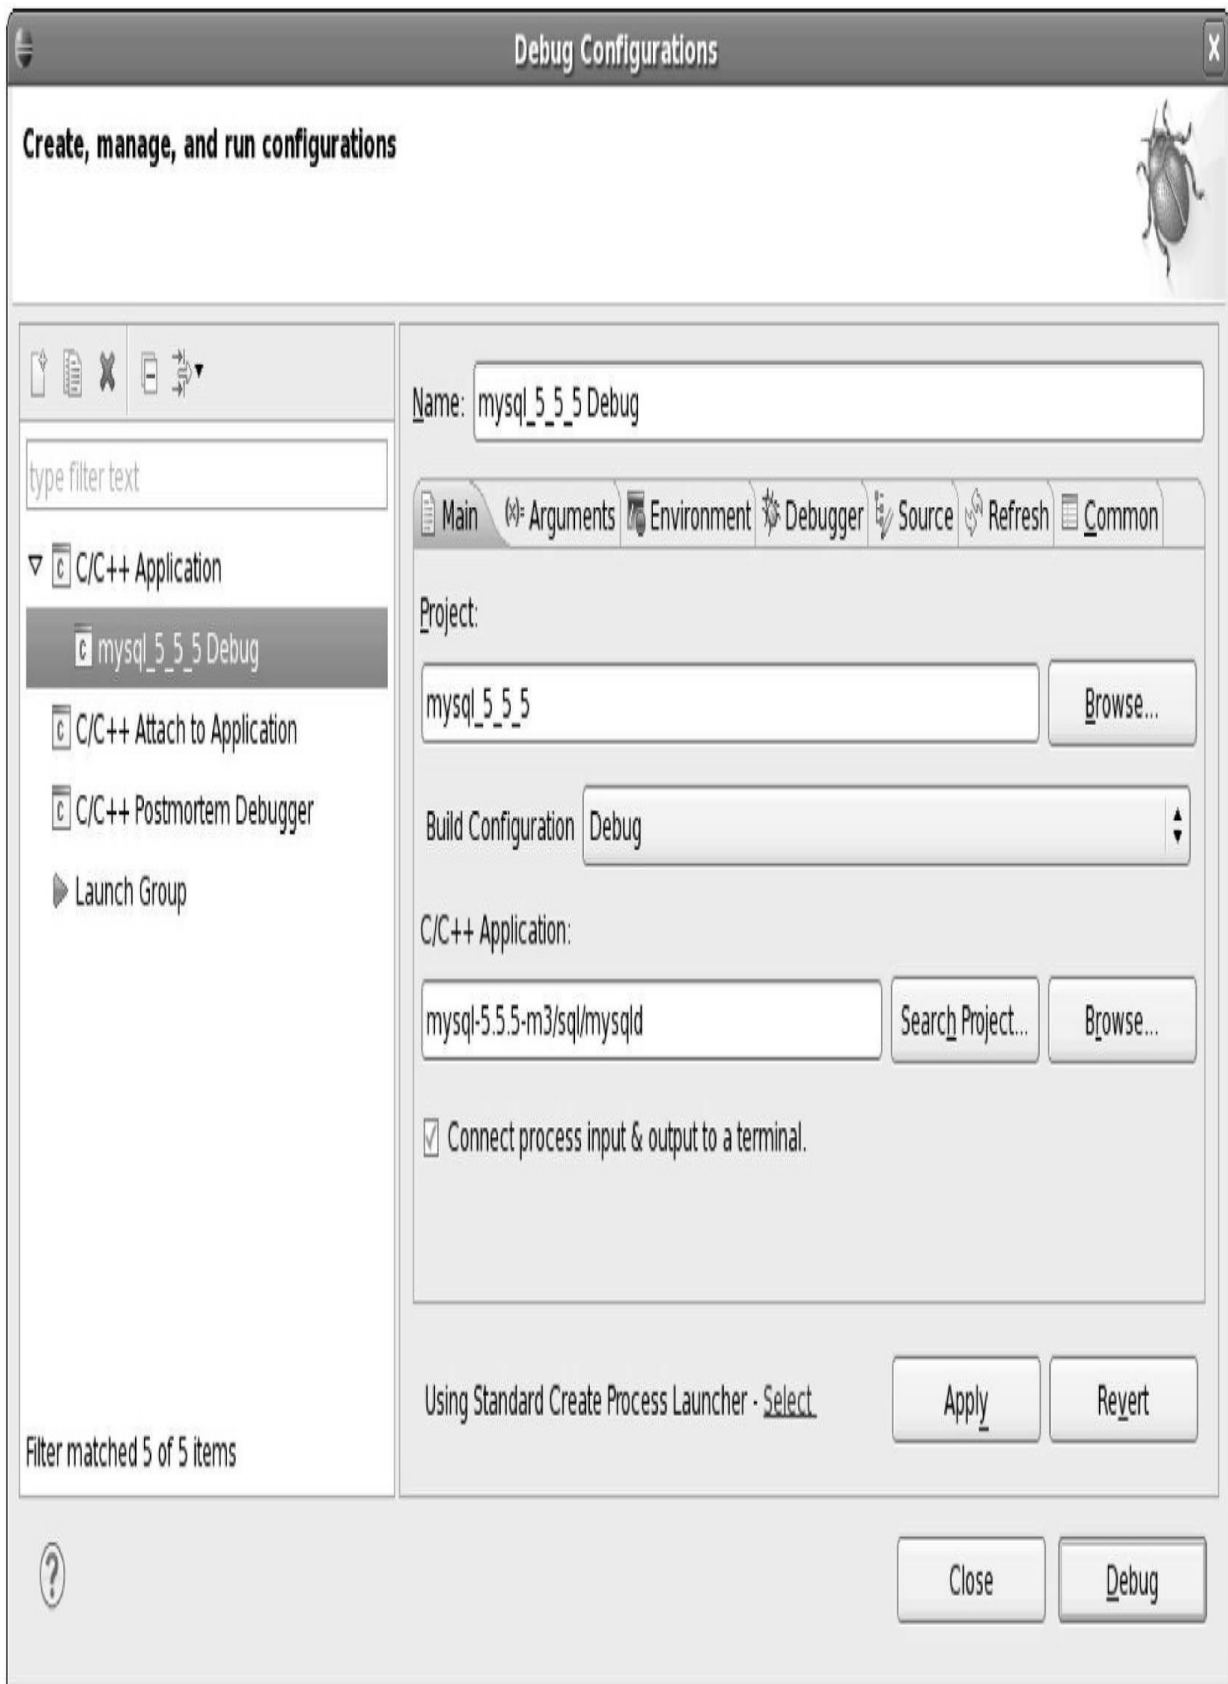
10-14

[illegible]



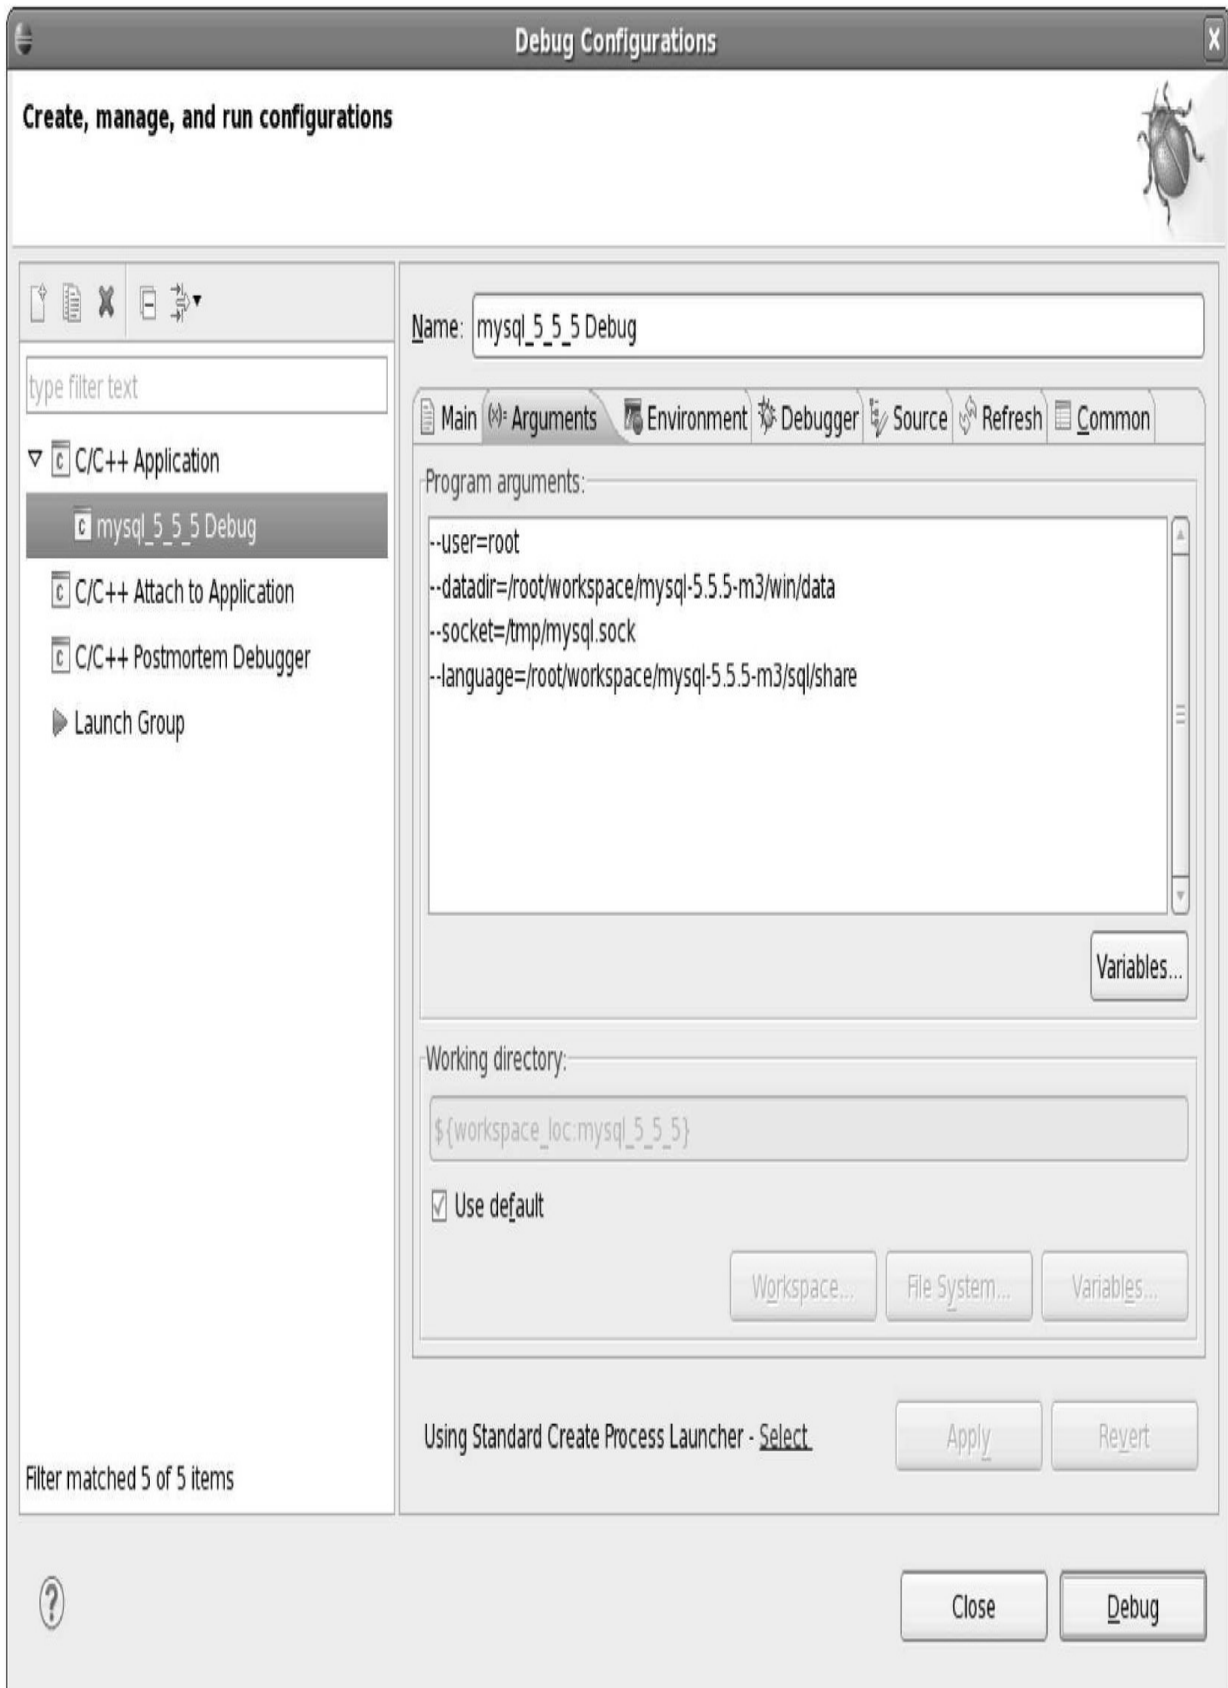
□ 10-15 □□□□

□□□□□□□□□□□□□□□□□□□□mysql□□□□□□□□□□□□□□□□
□□□□□Debug□□□□□□□□□□10-16□□□



□ 10-16 Debug□□

□□□□□□□□□□□□□□□□Arguments□□□□10-17□□□



10-17

10-18



□ 10-18 □Eclipse□□□□

10.4 cmake安装InnoDB

MySQL5.5cmakeMySQL
Mac OSXxcodemysql

```
cd mysql-xxx

mkdir bld

cd bld

cmake..-G Xcode
```

```

MySQLbldcmakexcode
MySQL.xcodeprojMySQL
InnoDB10-19

```


MySQL.xcodeproj — srv0srv.c

Running mysqld: ALL_BUILD

Project 257

Run Stop Scheme Breakpoints Editor View Organizer

ha_inn... sync0s... row0m... trx0trx.h ha_inn... row0sel.c btr0cur.c sync0s... srv0srv.c

MySQL > Sources > innobase > Source Files > srv0srv.c > srv_master_thread()

mysqld
Paused

- Thread 1
com.apple.main-thread
- Thread 2
 - 0 srv_master_thread
 - 1 pthread_start
 - 2 thread_start
 - Thread 3
 - Thread 4
 - Thread 5
 - Thread 6
 - Thread 7
 - Thread 8
 - Thread 9
 - Thread 10
 - Thread 11
 - Thread 12
 - Thread 13
 - Thread 14
 - Thread 15

```

2695 stot = srv_table_reserve_stot(SRV_MASTER);
2696
2697 srv_n_threads_active[SRV_MASTER]++;
2698
2699 mutex_exit(&kernel_mutex);
2700
2701 loop:
2702 /*-----*/
2703 /* --- When there is database activity by users, we cycle in this
2704 loop */
2705
2706 srv_main_thread_op_info = "reserving kernel mutex";
2707
2708 buf_get_total_stat(&buf_stat);
2709 n_ios_very_old = log_sys->n_log_ios + buf_stat.n_pages_read
2710 + buf_stat.n_pages_written;
2711 mutex_enter(&kernel_mutex);
2712
2713 /* Store the user activity counter at the start of this loop */
2714 old_activity_count = srv_activity_count;
2715
2716 mutex_exit(&kernel_mutex);
2717
2718 if (srv_force_recovery >= SRV_FORCE_NO_BACKGROUND) {
2719     goto suspend_thread;
2720 }
2721
2722 /* --- We run the following loop approximately once per second
2723 when there is database activity */
2724
2725 srv_last_log_flush_time = time(NULL);
2726
2727 /* Sleep for 1 second on entering the for loop below the first time. */
2728 next_itr_time = ut_time_ms() + 1000;
2729
2730 for (i = 0; i < 10; i++) {
2731

```

Thread 2: breakpoint 2.1

mysqld > Thread 2 > 0 srv_master_thread

Auto

- buf_stat (buf_pool_stat_t)
- old_activity_count = (uint) 0
- srv_activity_count = (uint) 0
- srv_force_recovery = (uint) 0
- block->lock = Invalid Expression

□ 10-19 Mac OS X□□□□□xcode□MySQL□□□□

□□cmake□□□□□MySQL□□□□□□□□cmake□□□□□□MySQL
□□□□□□

10.5 练习

MySQL、InnoDB、Visual Studio、Eclipse、Oracle、Microsoft SQL Server、DB2